

# Juniper Networks® Steel-Belted Radius® Carrier

---

## Installation Guide

Published  
2021-11-29

Release  
8.6.0

Juniper Networks, Inc.  
 1133 Innovation Way  
 Sunnyvale, California 94089  
 USA  
 408-745-2000  
[www.juniper.net](http://www.juniper.net)

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Ulticom, Signalware, Programmable Network, Ultimate Call Control, and Nexworx are registered trademarks of Ulticom, Inc. Kineto and the Kineto Logo are registered trademarks of Kineto Wireless, Inc. Software Advancing Communications and SignalCare are trademarks and service marks of Ulticom, Inc. CORBA (Common Object Request Broker Architecture) is a registered trademark of the Object Management Group (OMG). Raima, Raima Database Manager, and Raima Object Manager are trademarks of Raima, Inc. Sun, Sun Microsystems, the Sun logo, Java, Solaris, MySQL, and all trademarks and logos that contain Sun, Solaris, MySQL, or Java are trademarks or registered trademarks of Oracle America, Inc. in the United States and other countries. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. All specifications are subject to change without notice.

Contains software copyright 2000–2014 by Oracle America, Inc., distributed under license.

Steel-Belted Radius uses Thrift, licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the license at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

Steel-Belted Radius uses Xerces XML DOM, from the Apache Group. It has the following terms.

The Apache Software license, Version 1.1

Copyright © 1999-2003 The Apache Software Foundation. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The end-user documentation included with the redistribution, if any, must include the following acknowledgment: "This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>)." Alternately, this acknowledgment may appear in the software itself, if and wherever such third-party acknowledgments normally appear.
4. The names "Xerces" and "Apache Software Foundation" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact [apache@apache.org](mailto:apache@apache.org).
5. Products derived from this software may not be called "Apache", nor may "Apache" appear in their name, without prior written permission of the Apache Software Foundation.

THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the Apache Software Foundation and was originally based on software copyright © 1999, International Business Machines, Inc., <http://www.ibm.com>. For more information on the Apache Software Foundation, please see <http://www.apache.org/>.

Steel-Belted Radius uses the LDAP v2 Server from the University of Michigan. It has the following terms.

Copyright © 1991 Regents of the University of Michigan. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that this notice is preserved and that due credit is given to the University of Michigan at Ann Arbor. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission. This software is provided "as is" without express or implied warranty.

Portions of this software copyright 2003-2009 Lev Walkin <[vlm@lionet.info](mailto:vlm@lionet.info)> All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions of this software copyright 1989, 1991, 1992 by Carnegie Mellon University

SBR includes NetSNMP under the following licenses: Derivative Work-1996, 1998-2009 Copyright 1996, 1998-2009. The Regents of the University of California All Rights Reserved. Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Portions of this software copyright © 2001-2009, Networks Associates Technology, Inc. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the Networks Associates Technology, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR

OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions of this software are copyright © 2001–2009, Cambridge Broadband Ltd. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Steel-Belted Radius uses Jaxen, a “Java XPath Engine” from The Werken Company under the following license:

Copyright 2003 © The Werken Company. All Rights Reserved.

Redistribution and use of this software and associated documentation (“Software”), with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain copyright statements and notices. Redistributions must also contain a copy of this document.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name “jaxen” must not be used to endorse or promote products derived from this Software without prior written permission of The Werken Company. For written permission, please contact [bob@werken.com](mailto:bob@werken.com).
4. Products derived from this Software may not be called “jaxen” nor may “jaxen” appear in their names without prior written permission of The Werken Company. “jaxen” is a registered trademark of The Werken Company.
5. Due credit should be given to The Werken Company. (<http://jaxen.werken.com/>).

THIS SOFTWARE IS PROVIDED BY THE WERKEN COMPANY AND CONTRIBUTORS “AS IS” AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE WERKEN COMPANY OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR

CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

HTTPClient package Copyright © 1996–2009 Ronald Tschalär (ronald@innovation.ch)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details. For a copy of the GNU Lesser General Public License, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307, USA.

Steel-Belted Radius uses OpenSSL version 1.1.1, which have the following terms:

Copyright ©1998-2018 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit.  
(<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact [openssl-core@openssl.org](mailto:openssl-core@openssl.org).

5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.

6. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit  
(<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES

(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

OpenSSL is also subject to the following terms.

Copyright ©1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com).

The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed.

If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used.

This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)"

The word 'cryptographic' can be left out if the rouines from the library being used are not cryptographic related :

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT,

INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

SBR contains software copyright © 2000–2009 by The Legion Of The Bouncy Castle (<http://www.bouncycastle.org>)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Steel-Belted Radius uses modified source from OpenSolaris (now Oracle) under the CDDL, which can be found at [http://hub.opensolaris.org/bin/view/Main/opensolaris\\_license](http://hub.opensolaris.org/bin/view/Main/opensolaris_license). Modified source is available. Please refer to the SBR Carrier release notes.

SBR includes Spider Monkey libraries under Mozilla Public License Version 2.0

## 1. Definitions

1.1. “Contributor” means each individual or legal entity that creates, contributes to the creation of, or owns Covered Software.

1.2. “Contributor Version” means the combination of the Contributions of others (if any) used by a Contributor and that particular Contributor’s Contribution.

1.3. “Contribution” means Covered Software of a particular Contributor.

1.4. “Covered Software” means Source Code Form to which the initial Contributor has attached the notice in Exhibit A, the Executable Form of such Source Code Form, and Modifications of such Source Code Form, in each case including portions thereof.

1.5. “Incompatible With Secondary Licenses” means that the initial Contributor has attached the notice described in Exhibit B to the Covered Software; or that the Covered Software was made available under the terms of version 1.1 or earlier of the License, but not also under the terms of a Secondary License.

1.6. “Executable Form” means any form of the work other than Source Code Form.



1.7. “Larger Work” means a work that combines Covered Software with other material, in a separate file or files, that is not Covered Software.

1.8. “License” means this document.

1.9. “Licensable” means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently, any and all of the rights conveyed by this License.

1.10. “Modifications” means any of the following: any file in Source Code Form that results from an addition to, deletion from, or modification of the contents of Covered Software; or any new file in Source Code Form that contains any Covered Software.

1.11. “Patent Claims” of a Contributor means any patent claim(s), including without limitation, method, process, and apparatus claims, in any patent Licensable by such Contributor that would be infringed, but for the grant of the License, by the making, using, selling, offering for sale, having made, import, or transfer of either its Contributions or its Contributor Version.

1.12. “Secondary License” means either the GNU General Public License, Version 2.0, the GNU Lesser General Public License, Version 2.1, the GNU Affero General Public License, Version 3.0, or any later versions of those licenses.

1.13. “Source Code Form” means the form of the work preferred for making modifications.

1.14. “You” (or “Your”) means an individual or a legal entity exercising rights under this License. For legal entities, “You” includes any entity that controls, is controlled by, or is under common control with You. For purposes of this definition, “control” means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

## 2. License Grants and Conditions

### 2.1. Grants

Each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license: under intellectual property rights (other than patent or trademark) Licensable by such Contributor to use, reproduce, make available, modify, display, perform, distribute, and otherwise exploit its Contributions, either on an unmodified basis, with Modifications, or as part of a Larger Work; and under Patent Claims of such Contributor to make, use, sell, offer for sale, have made, import, and otherwise transfer either its Contributions or its Contributor Version.

### 2.2. Effective Date

The licenses granted in Section 2.1 with respect to any Contribution become effective for each Contribution on the date the Contributor first distributes such Contribution.

### 2.3. Limitations on Grant Scope

The licenses granted in this Section 2 are the only rights granted under this License. No additional rights or licenses will be implied from the distribution or licensing of Covered Software under this License. Notwithstanding Section 2.1(b) above, no patent license is granted by a Contributor: for any code that a Contributor has removed from Covered Software; or for infringements caused by: (i) Your and any other third party’s modifications of Covered Software, or (ii) the combination

of its Contributions with other software (except as part of its Contributor Version); or under Patent Claims infringed by Covered Software in the absence of its Contributions.

This License does not grant any rights in the trademarks, service marks, or logos of any Contributor (except as may be necessary to comply with the notice requirements in Section 3.4).

## 2.4. Subsequent Licenses

No Contributor makes additional grants as a result of Your choice to distribute the Covered Software under a subsequent version of this License (see Section 10.2) or under the terms of a Secondary License (if permitted under the terms of Section 3.3).

## 2.5. Representation

Each Contributor represents that the Contributor believes its Contributions are its original creation(s) or it has sufficient rights to grant the rights to its Contributions conveyed by this License.

## 2.6. Fair Use

This License is not intended to limit any rights You have under applicable copyright doctrines of fair use, fair dealing, or other equivalents.

## 2.7. Conditions

Sections 3.1, 3.2, 3.3, and 3.4 are conditions of the licenses granted in Section 2.1.

# 3. Responsibilities

## 3.1. Distribution of Source Form

All distribution of Covered Software in Source Code Form, including any Modifications that You create or to which You contribute, must be under the terms of this License. You must inform recipients that the Source Code Form of the Covered Software is governed by the terms of this License, and how they can obtain a copy of this License. You may not attempt to alter or restrict the recipients' rights in the Source Code Form.

## 3.2. Distribution of Executable Form

If You distribute Covered Software in Executable Form then:

such Covered Software must also be made available in Source Code Form, as described in Section 3.1, and You must inform recipients of the Executable Form how they can obtain a copy of such Source Code Form by reasonable means in a timely manner, at a charge no more than the cost of distribution to the recipient; and

You may distribute such Executable Form under the terms of this License, or sublicense it under different terms, provided that the license for the Executable Form does not attempt to limit or alter the recipients' rights in the Source Code Form under this License.

## 3.3. Distribution of a Larger Work

You may create and distribute a Larger Work under terms of Your choice, provided that You also comply with the requirements of this License for the Covered Software. If the Larger Work is a combination of Covered Software with a work governed by one or more Secondary Licenses, and the Covered Software is not Incompatible With Secondary Licenses, this License permits You to additionally distribute such Covered Software under the terms of such Secondary License(s), so that the recipient of the Larger Work may, at their option, further distribute the Covered Software under the terms of either this License or such Secondary License(s).

### 3.4. Notices

You may not remove or alter the substance of any license notices (including copyright notices, patent notices, disclaimers of warranty, or limitations of liability) contained within the Source Code Form of the Covered Software, except that You may alter any license notices to the extent required to remedy known factual inaccuracies.

### 3.5. Application of Additional Terms

You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, You may do so only on Your own behalf, and not on behalf of any Contributor. You must make it absolutely clear that any such warranty, support, indemnity, or liability obligation is offered by You alone, and You hereby agree to indemnify every Contributor for any liability incurred by such Contributor as a result of warranty, support, indemnity or liability terms You offer. You may include additional disclaimers of warranty and limitations of liability specific to any jurisdiction.

## 4. Inability to Comply Due to Statute or Regulation

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Software due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be placed in a text file included with all distributions of the Covered Software under this License. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

## 5. Termination

5.1. The rights granted under this License will terminate automatically if You fail to comply with any of its terms. However, if You become compliant, then the rights granted under this License from a particular Contributor are reinstated (a) provisionally, unless and until such Contributor explicitly and finally terminates Your grants, and (b) on an ongoing basis, if such Contributor fails to notify You of the non-compliance by some reasonable means prior to 60 days after You have come back into compliance. Moreover, Your grants from a particular Contributor are reinstated on an ongoing basis if such Contributor notifies You of the non-compliance by some reasonable means, this is the first time You have received notice of non-compliance with this License from such Contributor, and You become compliant prior to 30 days after Your receipt of the notice.

5.2. If You initiate litigation against any entity by asserting a patent infringement claim (excluding declaratory judgment actions, counter-claims, and cross-claims) alleging that a Contributor Version directly or indirectly infringes any patent, then the rights granted to You by any and all Contributors for the Covered Software under Section 2.1 of this License shall terminate.

5.3. In the event of termination under Sections 5.1 or 5.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or Your distributors under this License prior to termination shall survive termination.

## 6. Disclaimer of Warranty

Covered Software is provided under this License on an “as is” basis, without warranty of any kind, either expressed, implied, or statutory, including, without limitation, warranties that the Covered Software is free of defects, merchantable, fit for a particular purpose or non-infringing. The entire risk as to the quality and performance of the Covered Software is with You. Should any Covered Software prove defective in any respect, You (not any Contributor) assume the cost of any necessary servicing, repair, or correction. This disclaimer of warranty constitutes an essential part of this License. No use of any Covered Software is authorized under this License except under this disclaimer.

## 7. Limitation of Liability

Under no circumstances and under no legal theory, whether tort (including negligence), contract, or otherwise, shall any Contributor, or anyone who distributes Covered Software as permitted above, be liable to You for any direct, indirect, special, incidental, or consequential damages of any character including, without limitation, damages for lost profits, loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses, even if such party shall have been informed of the possibility of such damages. This limitation of liability shall not apply to liability for death or personal injury resulting from such party's negligence to the extent applicable law prohibits such limitation. Some jurisdictions do not allow the exclusion or limitation of incidental or consequential damages, so this exclusion and limitation may not apply to You.

## 8. Litigation

Any litigation relating to this License may be brought only in the courts of a jurisdiction where the defendant maintains its principal place of business and such litigation shall be governed by laws of that jurisdiction, without reference to its conflict-of-law provisions. Nothing in this Section shall prevent a party's ability to bring cross-claims or counter-claims.

## 9. Miscellaneous

This License represents the complete agreement concerning the subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not be used to construe this License against a Contributor.

## 10. Versions of the License

### 10.1. New Versions

Mozilla Foundation is the license steward. Except as provided in Section 10.3, no one other than the license steward has the right to modify or publish new versions of this License. Each version will be given a distinguishing version number.

### 10.2. Effect of New Versions

You may distribute the Covered Software under the terms of the version of the License under which You originally received the Covered Software, or under the terms of any subsequent version published by the license steward.

### 10.3. Modified Versions

If you create software not governed by this License, and you want to create a new license for such software, you may create and use a modified version of this License if you rename the license and remove any references to the name of the license steward (except to note that such modified license differs from this License).

### 10.4. Distributing Source Code Form that is Incompatible With Secondary Licenses

If You choose to distribute Source Code Form that is Incompatible With Secondary Licenses under the terms of this version of the License, the notice described in Exhibit B of this License must be attached.

#### Exhibit A - Source Code Form License Notice

This Source Code Form is subject to the terms of the Mozilla Public License v.2.0. If a copy of the MPL was not distributed with this file, You can obtain one at <http://mozilla.org/MPL/2.0/>.

If it is not possible or desirable to put the notice in a particular file, then You may include the notice in a location (such as a LICENSE file in a relevant directory) where a recipient would be likely to look for such a notice.

You may add additional accurate notices of copyright ownership.

#### Exhibit B - "Incompatible With Secondary Licenses" Notice

This Source Code Form is "Incompatible With Secondary Licenses", as defined by the Mozilla Public License, v. 2.0.

The "inih" library is distributed under the New BSD license:

Copyright © 2009, Brush Technology  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of Brush Technology nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY BRUSH TECHNOLOGY "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL BRUSH TECHNOLOGY BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Contains software copyright 2007-2014, by Sencha, Inc., distributed under license.

Steel-Belted Radius uses Jetty 9 under the Apache License 2.0, You may obtain a copy of the license at <http://www.apache.org/licenses/LICENSE-2.0>

Steel-Belted Radius uses Google Web Toolkit (GWT) under the Apache License 2.0, You may obtain a copy of the license at <http://www.apache.org/licenses/LICENSE-2.0>

Steel-Belted Radius uses Apache HTTP components under the Apache License 2.0, You may obtain a copy of the license at <http://www.apache.org/licenses/LICENSE-2.0>

Steel-Belted Radius uses OpenJDK

GNU General Public License, version 2, with the Classpath Exception

The GNU General Public License (GPL)

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.  
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

## Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

## TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program

itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this



License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

## NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

## END OF TERMS AND CONDITIONS

### How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

One line to give the program's name and a brief idea of what it does.

Copyright © <year> <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright © year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type 'show w'. This is free software, and you are welcome to redistribute it under certain conditions; type 'show c' for details.

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than 'show w' and 'show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program 'Gnomovision' (which makes passes at compilers) written by James Hacker.  
signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

#### "CLASSPATH" EXCEPTION TO THE GPL

Certain source files distributed by Oracle America and/or its affiliates are subject to the following clarification and special exception to the GPL, but only where Oracle has expressly included in the particular source file's header the words "Oracle designates this particular file as subject to the "Classpath" exception as provided by Oracle in the LICENSE file that accompanied this code."

Linking this library statically or dynamically with other modules is making a combined work based on this library. Thus, the terms and conditions of the GNU General Public License cover the whole combination.

As a special exception, the copyright holders of this library give you permission to link this library with independent modules to produce an executable, regardless of the license terms of these independent modules, and to copy and distribute the resulting executable under terms of your choice, provided that you also meet, for each linked independent module, the terms and conditions of the license of that module. An independent module is a module which is not derived from or based on this library. If you modify this library, you may extend this exception to your version of the library, but you are not obligated to do so. If you do not wish to do so, delete this exception statement from your version.

SBR uses Gecko SDK 1.4b

Mozilla Public License Version 2.0

#### 1. Definitions

1.1. "Contributor" means each individual or legal entity that creates, contributes to the creation of, or owns Covered Software.

1.2. "Contributor Version" means the combination of the Contributions of others (if any) used by a Contributor and that particular Contributor's Contribution.

1.3. "Contribution" means Covered Software of a particular Contributor.

1.4. "Covered Software" means Source Code Form to which the initial Contributor has attached the notice in Exhibit A, the Executable Form of such Source Code Form, and Modifications of such Source Code Form, in each case including portions thereof.

1.5. "Incompatible With Secondary Licenses" means

(a) that the initial Contributor has attached the notice described in Exhibit B to the Covered Software; or

(b) that the Covered Software was made available under the terms of version 1.1 or earlier of the License, but not also under the terms of a Secondary License.

1.6. "Executable Form" means any form of the work other than Source Code Form.

1.7. "Larger Work" means a work that combines Covered Software with other material, in a separate file or files, that is not Covered Software.

1.8. "License" means this document.

1.9. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently, any and all of the rights conveyed by this License.

1.10. "Modifications" means any of the following:

(a) any file in Source Code Form that results from an addition to, deletion from, or modification of the contents of Covered Software; or

(b) any new file in Source Code Form that contains any Covered Software.

1.11. "Patent Claims" of a Contributor means any patent claim(s), including without limitation, method, process, and apparatus claims, in any patent Licensable by such Contributor that would be infringed, but for the grant of the License, by the making, using, selling, offering for sale, having made, import, or transfer of either its Contributions or its Contributor Version.

1.12. "Secondary License" means either the GNU General Public License, Version 2.0, the GNU Lesser General Public License, Version 2.1, the GNU Affero General Public License, Version 3.0, or any later versions of those licenses.

1.13. "Source Code Form" means the form of the work preferred for making modifications.

1.14. "You" (or "Your") means an individual or a legal entity exercising rights under this License. For legal entities, "You" includes any entity that controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

## 2. License Grants and Conditions

### 2.1. Grants

Each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license:

(a) under intellectual property rights (other than patent or trademark) Licensable by such Contributor to use, reproduce, make available, modify, display, perform, distribute, and otherwise exploit its Contributions, either on an unmodified basis, with Modifications, or as part of a Larger Work; and

(b) under Patent Claims of such Contributor to make, use, sell, offer for sale, have made, import, and otherwise transfer either its Contributions or its Contributor Version.

## 2.2. Effective Date

The licenses granted in Section 2.1 with respect to any Contribution become effective for each Contribution on the date the Contributor first distributes such Contribution.

## 2.3. Limitations on Grant Scope

The licenses granted in this Section 2 are the only rights granted under this License. No additional rights or licenses will be implied from the distribution or licensing of Covered Software under this License. Notwithstanding Section 2.1(b) above, no patent license is granted by a Contributor:

(a) for any code that a Contributor has removed from Covered Software; or

b) for infringements caused by: (i) Your and any other third party's modifications of Covered Software, or (ii) the combination of its Contributions with other software (except as part of its Contributor Version); or

(c) under Patent Claims infringed by Covered Software in the absence of its Contributions.

This License does not grant any rights in the trademarks, service marks, or logos of any Contributor (except as may be necessary to comply with the notice requirements in Section 3.4).

## 2.4. Subsequent Licenses

No Contributor makes additional grants as a result of Your choice to distribute the Covered Software under a subsequent version of this License (see Section 10.2) or under the terms of a Secondary License (if permitted under the terms of Section 3.3).

## 2.5. Representation

Each Contributor represents that the Contributor believes its Contributions are its original creation(s) or it has sufficient rights to grant the rights to its Contributions conveyed by this License.

## 2.6. Fair Use

This License is not intended to limit any rights You have under applicable copyright doctrines of fair use, fair dealing, or other equivalents.

## 2.7. Conditions

Sections 3.1, 3.2, 3.3, and 3.4 are conditions of the licenses granted in Section 2.1.

### 3. Responsibilities

#### 3.1. Distribution of Source Form

All distribution of Covered Software in Source Code Form, including any Modifications that You create or to which You contribute, must be under the terms of this License. You must inform recipients that the Source Code Form of the Covered Software is governed by the terms of this License, and how they can obtain a copy of this License. You may not attempt to alter or restrict the recipients' rights in the Source Code Form.

#### 3.2. Distribution of Executable Form

If You distribute Covered Software in Executable Form then:

(a) such Covered Software must also be made available in Source Code Form, as described in Section 3.1, and You must inform recipients of the Executable Form how they can obtain a copy of such Source Code Form by reasonable means in a timely manner, at a charge no more than the cost of distribution to the recipient; and

(b) You may distribute such Executable Form under the terms of this License, or sublicense it under different terms, provided that the license for the Executable Form does not attempt to limit or alter the recipients' rights in the Source Code Form under this License.

#### 3.3. Distribution of a Larger Work

You may create and distribute a Larger Work under terms of Your choice, provided that You also comply with the requirements of this License for the Covered Software. If the Larger Work is a combination of Covered Software with a work governed by one or more Secondary Licenses, and the Covered Software is not Incompatible With Secondary Licenses, this License permits You to additionally distribute such Covered Software under the terms of such Secondary License(s), so that the recipient of the Larger Work may, at their option, further distribute the Covered Software under the terms of either this License or such Secondary License(s).

#### 3.4. Notices

You may not remove or alter the substance of any license notices (including copyright notices, patent notices, disclaimers of warranty, or limitations of liability) contained within the Source Code Form of the Covered Software, except that You may alter any license notices to the extent required to remedy known factual inaccuracies.

#### 3.5. Application of Additional Terms

You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, You may do so only on Your own behalf, and not on behalf of any Contributor. You must make it absolutely clear that any such warranty, support, indemnity, or liability obligation is offered by You alone, and You hereby agree to indemnify every Contributor for any liability incurred by such Contributor as a result of warranty, support, indemnity or liability terms You offer. You may include additional disclaimers of warranty and limitations of liability specific to any jurisdiction.

### 4. Inability to Comply Due to Statute or Regulation

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Software due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the

maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be placed in a text file included with all distributions of the Covered Software under this License. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

## 5. Termination

5.1. The rights granted under this License will terminate automatically if You fail to comply with any of its terms. However, if You become compliant, then the rights granted under this License from a particular Contributor are reinstated (a) provisionally, unless and until such Contributor explicitly and finally terminates Your grants, and (b) on an ongoing basis, if such Contributor fails to notify You of the non-compliance by some reasonable means prior to 60 days after You have come back into compliance. Moreover, Your grants from a particular Contributor are reinstated on an ongoing basis if such Contributor notifies You of the non-compliance by some reasonable means, this is the first time You have received notice of non-compliance with this License from such Contributor, and You become compliant prior to 30 days after Your receipt of the notice.

5.2. If You initiate litigation against any entity by asserting a patent infringement claim (excluding declaratory judgment actions, counter-claims, and cross-claims) alleging that a Contributor Version directly or indirectly infringes any patent, then the rights granted to You by any and all Contributors for the Covered Software under Section 2.1 of this License shall terminate.

5.3. In the event of termination under Sections 5.1 or 5.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or Your distributors under this License prior to termination shall survive termination.

## 6. Disclaimer of Warranty

Covered Software is provided under this License on an "as is" basis, without warranty of any kind, either expressed, implied, or statutory, including, without limitation, warranties that the Covered Software is free of defects, merchantable, fit for a particular purpose or non-infringing. The entire risk as to the quality and performance of the Covered Software is with You. Should any Covered Software prove defective in any respect, You (not any Contributor) assume the cost of any necessary servicing, repair, or correction. This disclaimer of warranty constitutes an essential part of this License. No use of any Covered Software is authorized under this License except under this disclaimer.

## 7. Limitation of Liability

Under no circumstances and under no legal theory, whether tort (including negligence), contract, or otherwise, shall any Contributor, or anyone who distributes Covered Software as permitted above, be liable to You for any direct, indirect, special, incidental, or consequential damages of any character including, without limitation, damages for lost profits, loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses, even if such party shall have been informed of the possibility of such damages. This limitation of liability shall not apply to liability for death or personal injury resulting from such party's negligence to the extent applicable law prohibits such limitation. Some jurisdictions do not allow the exclusion or limitation of incidental or consequential damages, so this exclusion and limitation may not apply to You.

## 8. Litigation

Any litigation relating to this License may be brought only in the courts of a jurisdiction where the defendant maintains its principal place of business and such litigation shall be governed by laws of that jurisdiction, without reference to its conflict-of-law provisions. Nothing in this Section shall prevent a party's ability to bring cross-claims or counter-claims.

## 9. Miscellaneous

This License represents the complete agreement concerning the subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not be used to construe this License against a Contributor.

## 10. Versions of the License

### 10.1. New Versions

Mozilla Foundation is the license steward. Except as provided in Section 10.3, no one other than the license steward has the right to modify or publish new versions of this License. Each version will be given a distinguishing version number.

### 10.2. Effect of New Versions

You may distribute the Covered Software under the terms of the version of the License under which You originally received the Covered Software, or under the terms of any subsequent version published by the license steward.

### 10.3. Modified Versions

If you create software not governed by this License, and you want to create a new license for such software, you may create and use a modified version of this License if you rename the license and remove any references to the name of the license steward (except to note that such modified license differs from this License).

### 10.4. Distributing Source Code Form that is Incompatible With Secondary Licenses

If You choose to distribute Source Code Form that is Incompatible With Secondary Licenses under the terms of this version of the License, the notice described in Exhibit B of this License must be attached.

#### Exhibit A - Source Code Form License Notice

This Source Code Form is subject to the terms of the Mozilla Public License, v. 2.0. If a copy of the MPL was not distributed with this file, You can obtain one at <http://mozilla.org/MPL/2.0/>.

If it is not possible or desirable to put the notice in a particular file, then You may include the notice in a location (such as a LICENSE file in a relevant directory) where a recipient would be likely to look for such a notice.

You may add additional accurate notices of copyright ownership.

#### Exhibit B - "Incompatible With Secondary Licenses" Notice

This Source Code Form is "Incompatible With Secondary Licenses", as defined by the Mozilla Public License, v. 2.0.

SBR uses Mozilla LDAP C SDK 5.17



## MOZILLA PUBLIC LICENSE

### Version 1.1

#### 1. Definitions

1.0.1. "Commercial Use" means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor" means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code" means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism" means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable" means Covered Code in any form other than Source Code.

1.6. "Initial Developer" means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.7. "Larger Work" means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License" means this document.

1.8.1. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

A. Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

B. Any new file that contains any part of the Original Code or previous Modifications.

1.10. "Original Code" means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code" means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code

of the Contributor's choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.12. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

## 2. Source Code License.

### 2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

### 2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version; 3) for infringements caused by: i) third party

modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

### 3. Distribution Obligations.

#### 3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

#### 3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

#### 3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

#### 3.4. Intellectual Property Matters

(a) Third Party Claims. If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

(b) Contributor APIs. If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations. Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

### 3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear than any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.6. Distribution of Executable Versions. You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligations of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7. Larger Works. You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

### 4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

### 5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

## 6. Versions of the License.

6.1. New Versions. Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2. Effect of New Versions. Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by Netscape. No one other than Netscape has the right to modify the terms applicable to Covered Code created under this License.

6.3. Derivative Works. If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the Mozilla Public License and Netscape Public License. (Filling in the name of the Initial Developer, Original Code or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

## 7. DISCLAIMER OF WARRANTY.

COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

## 8. TERMINATION.

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2. If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as "Participant") alleging that:

(a) such Participant's Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

(b) any software, hardware, or device, other than such Participant's Contributor Version, directly or indirectly infringes any patent, then any rights granted to You by such Participant under Sections 2.1(b) and 2.2(b) are revoked effective as of the date You first made, used, sold, distributed, or had made, Modifications made by that Participant.

8.3. If You assert a patent infringement claim against Participant alleging that such Participant's Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

## 9. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

## 10. U.S. GOVERNMENT END USERS.

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.

## 11. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. With respect to disputes in which at least one party is a citizen of, or an entity chartered or registered to do business in the United States of America, any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California, with venue lying in Santa Clara County, California, with the losing party responsible for costs, including without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

## 12. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

### 13. MULTIPLE-LICENSED CODE.

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the MPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

#### EXHIBIT A -Mozilla Public License.

"The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/MPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is \_\_\_\_\_.

The Initial Developer of the Original Code is \_\_\_\_\_. Portions created by \_\_\_\_\_ are Copyright (C) \_\_\_\_\_. All Rights Reserved.

Contributor(s): \_\_\_\_\_.

Alternatively, the contents of this file may be used under the terms of the \_\_\_\_\_ license (the "[\_\_\_\_\_] License"), in which case the provisions of [\_\_\_\_\_] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [\_\_\_\_\_] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [\_\_\_\_\_] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [\_\_\_\_\_] License."

[NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.]

*Steel-Belted Radius Carrier 8.6.0 Installation Guide*  
Release 8.6.0

#### Revision History

August 2019—Revision 1

October 2019—Revision 2

March 2020—Revision 3

September 2020—Revision 4

The information in this document is current as of the date on the title page.

#### YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

## END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.



# Abbreviated Table of Contents

**About This Guide | xlvii**

1

**Overview of Steel-Belted Radius Carrier and Session State Register Installation**

**Chapter 1 Standalone Steel-Belted Radius Carrier Installation Overview | 2**

**Chapter 2 Session State Register Cluster Installation Overview | 9**

2

**Preparing for a Steel-Belted Radius Carrier Installation**

**Chapter 3 Planning Your Session State Register Cluster | 32**

**Chapter 4 Before You Install Software | 39**

**Chapter 5 Migrating from Previous SBR Releases | 56**

3

**Installing and Configuring a SBR Carrier Standalone Server**

**Chapter 6 Installation and Basic Configuration of a SBR Carrier Standalone Server | 64**

4

## Installing Session State Register Nodes

**Chapter 7**    **Installing Session State Register Nodes | 94**

**Chapter 8**    **Basic SBR Carrier Node Configuration | 138**

**Chapter 9**    **Customizing the SSR Database Current Sessions Table | 144**

**Chapter 10**   **Other Session State Register Configuration Files | 182**

**Chapter 11**   **When and How to Restart Session State Register Nodes, Hosts, and Clusters | 196**

5

## Expanding a Session State Register Cluster

**Chapter 12**   **Expanding an SSR Cluster | 207**

6

## SIGTRAN Support for Steel-Belted Radius Carrier

**Chapter 13**   **Migrating Signalware | 310**

**Chapter 14**   **Installing Signalware 9 on Linux | 315**

**Chapter 15**   **Configuring SIGTRAN/IP Network Communication Files | 332**

7

Uninstalling Steel-Belted Radius Carrier Software

**Chapter 16 Removing Steel-Belted Radius Carrier Software | 359**

8

Upgrading Your SSR Cluster

**Chapter 17 Overview of Upgrading Your SSR Cluster | 366****Chapter 18 Using a Transition Server to Mitigate Downtime While Upgrading Your Cluster | 369****Chapter 19 Upgrading Your Cluster Using the Rolling Restart Method | 396****Chapter 20 Upgrading Your Cluster Using the Backup, Destroy, and Re-Create Method | 447****Part 9 Appendix****Appendix A Testing the Installation with DemoSetup.sh | 482**

# Table of Contents

## About This Guide | xlvii

Objective | xlvii

Audience | xlvii

Documentation Conventions | xlviii

Related Documentation | i

Obtaining Documentation | lvi

Documentation Feedback | lvi

Requesting Technical Support | lvii

1

## Overview of Steel-Belted Radius Carrier and Session State Register Installation

### Standalone Steel-Belted Radius Carrier Installation Overview | 2

Migrating from Earlier SBR Releases | 2

Supporting SIM with Signalware | 3

Installation Workflow | 4

Release 8.6.0 Standalone Server Installation Checklist | 4

Before You Begin | 5

Installing SBR Software | 7

Basic Configuration | 7

### Session State Register Cluster Installation Overview | 9

SSR Cluster Overview | 9

SSR Cluster Concepts and Terminology | 10

Session State Register Servers | 10

Session State Register Nodes | 11

SSR Data Entities | 11

Cluster Configurations | 13

Session State Register Scaling | 13

Adding a Data Node Expansion Kit | 14

Adding a Third Management Node | 14

	Adding More SBR Carrier Front End Servers   14
	Cluster Network Requirements   15
	Supported SBR Carrier SSR Cluster Configurations   16
	Failover Overview   17
	Failover Examples   18
	Distributed Cluster Failure and Recovery   20
	Session State Register Database Tables   22
	IP Address Pools   22
	Subscriber Session Data Controls   23
	Application Support   24
	Supporting SIM, SIGTRAN, and Other Protocols with Signalware   25
	Cluster Installation Workflow   25
	Session State Register Cluster Installation Checklist   26
	Before You Begin   26
	Installing Session State Register Software   29
	Basic Configuration   29

## 2

## Preparing for a Steel-Belted Radius Carrier Installation

### Planning Your Session State Register Cluster | 32

#### Planning the Cluster Topology | 32

##### Naming the Cluster and Its Machines | 32

##### Starter Kit Cluster Naming Example | 33

#### Cluster Naming Worksheets | 34

##### Expansion Kit and Additional Node Worksheet | 35

#### Renaming Node Hosts | 38

#### Configuring Multipathing | 38

### Before You Install Software | 39

#### Reviewing the Release Notes | 39

#### Determining the Server's Centralized Configuration Management Role | 40

#### Selecting an Appropriate Server | 41

**Meeting System Requirements | 41****Standalone SBR Carrier Server Hardware | 41****Session State Register Host Hardware | 43****SBR Carrier and Management Node Hosts | 43****Data Node Hosts | 44****Checking Free Disk Space | 45****Software | 46****Solaris | 47****Linux | 47****Perl | 48****LDAP Plug-in | 49****Verifying Root Access | 49****Verifying Network Requirements | 49****Verifying Network Connectivity | 49****Verifying Hostname Resolution | 50****Creating Required Users and Groups | 50****Creating Share Directories | 52****Setting Up External Database Connectivity (Optional) | 53****Creating a Copy of Existing SBR Server Release Files for Migration | 54****Installing the SIGTRAN Interface (Optional) | 55****Migrating from Previous SBR Releases | 56****Migration Overview | 56****Supported Releases | 57****Files That May Be Migrated | 57****Files from Earlier Releases That Require Manual Editing | 58****Manual Migration of XML Configurations | 58****Manual Migration of ROOT Certificates | 58****Manual Migration of Dictionaries | 59****Manual Migration of Third-Party Plug-ins and Other Binaries | 59****Manual Migration of SNMP Files | 59****Migrating from SBR Release 6.X | 59****Migrating from SBR Carrier Release 7.0 | 59****Migrating from SBR Carrier Release 7.4.0 | 60**

Migration and New Installations of SBR Carrier with WiMAX | 60

Historic Files | 61

### 3

## Installing and Configuring a SBR Carrier Standalone Server

### Installation and Basic Configuration of a SBR Carrier Standalone Server | 64

Unpacking the Steel-Belted Radius Carrier Software | 64

Package Management Commands | 64

Running the Steel-Belted Radius Carrier Configure Script | 68

Starting and Stopping a Standalone Steel-Belted Radius Carrier Server | 77

sbrd | 77

Running sbrd on a Standalone Server | 77

Starting the RADIUS Server | 78

Stopping the RADIUS Server | 79

Displaying RADIUS Status Information | 79

Basic SBR Carrier Configuration | 84

Launching Web GUI | 84

Configuring the Server | 86

Configuring SNMP | 86

Configuring the CST for a Standalone Server | 88

Configuring the sessionTable.ini file | 89

RadAttr Fields and the sessionTable.ini File | 89

Multi-Valued Attributes | 90

Using MVA Facilities | 92

### 4

## Installing Session State Register Nodes

### Installing Session State Register Nodes | 94

Cluster Installation Overview | 94

Unpacking Session State Register Software | 95

Package Management Commands | 95

Setting Up a Starter Kit's First SBR/Management Node | 99

Configuring the Host Software on the First Server in the Cluster | 99

Setting Up the Second SBR/Management Node in a Starter Kit | 114

Populating the JNPRShare Directory | 115

Configuring Host Software on the Second SBR Carrier and Management Node | 115

## Setting Up Data Node Hosts Included with the Starter Kit | 126

- Striping Data Nodes | 126

- Populating the JNPRShare Directory | 128

- Configuring the Host Software on the Data Nodes | 129

## Initial Cluster Startup Procedure | 132

## Recommendations before Configuring the Cluster | 137

## Basic SBR Carrier Node Configuration | 138

- Launching Web GUI | 138

- Configuring the Server | 140

- Configuring SNMP | 141

- Setting Up IP Address Pools | 142

## Customizing the SSR Database Current Sessions Table | 144

### Current Sessions Table Overview | 144

### Sessions in the CST | 153

- Acknowledging Authentication and Accounting Requests when the CST Cannot Be Contacted | 153

- Current Sessions Table Display | 154

### Customizing the CST | 154

- Propagating a Changed CurrentSessions.sql File | 155

- Performance and Capacity Considerations | 155

- Additional Keys | 155

- Stored Procedures | 155

- Customized CST Applications | 156

### Current Sessions Table Fields | 156

- System Core Fields | 156

- System Feature Fields | 159

- System Optional Fields | 160

- RadAttr Fields | 162

- RadAttr Fields and the sessionTable.ini File | 164

- Multi-Valued Attributes | 167

- Using MVA Facilities | 169

- Admin Private Fields | 170

- System Keys and Indexes | 170



Admin Keys and Indexes | 171

SSR Datatypes | 172

Definitions | 172

System Field Datatypes | 173

RadAttr Fields | 174

Matching/Natural Datatype Mappings | 174

Sample ShowSessions.sh Report | 175

Compatible/Unnatural Datatype Conversions | 178

Customizing Datatypes and Data Sizes | 179

Character Sets and Collations | 180

Truncation and Padding | 180

Small Integers | 181

Signed Integers | 181

Unsupported SQL Datatypes | 181

**Other Session State Register Configuration Files | 182**

Tuning Other SSR Tables Overview | 182

Configuring the dbclusterndb.gen File | 182

[Bootstrap] Section | 182

[NDB] Section | 183

[Database] Section | 186

[IpAddressPools] Section | 188

[IpAddressPools:PoolName] Section | 191

Customizing Other Tables and Scripts | 193

Managing IP Pools | 194

**When and How to Restart Session State Register Nodes, Hosts, and Clusters | 196**

Overview of Starting and Stopping a Session State Register Cluster | 196

sbrd | 197

Running sbrd on Session State Register Nodes | 197

Starting the Cluster | 200

Proper Order for Starting Nodes in a Cluster | 201

Proper Order for Stopping Individual Nodes | 203

Stopping a Single Node | 204

Starting a Single Node | 204

Proper Order for Stopping the Entire Cluster | 205

## Expanding a Session State Register Cluster

### Expanding an SSR Cluster | 207

Upgrading from a Restricted Cluster to a Regular Cluster | 207

Adding Nodes to a Cluster Overview | 210

Adding a New SBR Carrier Server to an Existing Cluster | 211

Updating the Existing Cluster Definition Files for the New SBR Node | 212

Distributing the Updated Cluster Definition Files to the Existing Nodes | 218

Installing the SBR Carrier Software on the New SBR Node Host Machine | 220

Configuring the Software on the New SBR Node | 222

Configuring Each Existing Node in the Cluster with the New Cluster Definition Files | 228

Starting the New SBR Node | 235

Adding a Management Node Expansion Kit to an Existing Cluster | 235

Updating the Existing Cluster Definition Files for the New Management Node | 237

Distributing the Updated Cluster Definition Files to the Existing Nodes | 244

Installing the SBR Carrier Software on the New Management Node Host Machine | 245

Configuring the SBR Carrier Software on the New Management Node | 248

Configuring Each Existing Node in the Cluster with the New Cluster Definition Files | 252

Starting the New Management Node | 259

Running CreateDB.sh on the New Management Node | 259

Adding a Data Expansion Kit to an Existing Cluster | 259

Requirements for Selecting a Transition Server in Your Environment | 260

Using a Transition Server When Adding Data Nodes to an Existing Cluster | 261

Existing Cluster Configuration for This Example Procedure | 261

Creating the Transition Server | 263

Stopping the Processes on the Target Transition Server | 263

Configuring the Software on the Transition Server as a Temporary Cluster | 264

Configuring and Starting the Transition Server | 270

Switching Traffic to the Transition Server | 271

Creating the Updated Cluster Definition Files | 271

Installing the SBR Carrier Software on the Two New Data Node Host Machines | 278

Distributing the Updated Cluster Definition Files to the Existing Nodes | 281

Destroying the Session Database on the Original Cluster	282
Configuring the Nodes in the Expanded Cluster with the Updated Cluster Definition Files	284
Configuring the SBR Carrier Software on the New Data Nodes	284
Running the Configure Script on Each Node from the Original Cluster	289
Creating the Session Database and IP Pools on the Expanded Cluster	293
Cleaning the Original Nodes from the Cluster	293
Creating the Session Database and IP Pools	294
Removing the Transition Server from Service	297
Unconfiguring and Rebuilding the Transition Server	297
Unconfiguring the Transition Server	297
Retrieving the Updated Cluster Definition Files from SM1 Node	300
Running the Configure Script on the SM2 Node	300
Non-Transition Server Method—Terminating Connections	307

## **SIGTRAN Support for Steel-Belted Radius Carrier**

### **Migrating Signalware | 310**

Migrating an Existing Signalware Configuration	310
Backing Up the Existing Signalware Configuration	310
Applying the Saved Signalware Configuration to the New Server	312

### **Installing Signalware 9 on Linux | 315**

Operating System Prerequisite Packages	316
64-Bit x86_64 Red Hat Installations	316
Red Hat Enterprise Linux Release 7.3 and Above	317
Other Prerequisites to Install Signalware on a Linux Platform	318
Installing Signalware	319
Troubleshooting Signalware Installation and Configuration	328

### **Configuring SIGTRAN/IP Network Communication Files | 332**

Starting and Stopping Signalware	332
Configuring Signalware to Start Automatically on Reboot	332
Configuring Communication Files Overview	333
Communication Pathways and Corresponding Files	334
Configuration Activities	335
Signalware MML Commands	335

Defining Links, Link Sets, and Route Sets | 336

Example MML Commands | 336

SS7 Example: Creating Links, Link Sets, and Route Sets | 337

SIGTRAN Example: Creating Links, Link Sets, and Route Sets | 337

Configuring authGateway and GWrelay Applications for HLR Communication | 338

Configuring the authGateway Routing Location Information | 338

Example 1—Global Title Routing Using Global Title Identification | 339

Example 2—PCSSN Routing Using Point Code Identification | 339

Configuring the authGateway.conf File | 339

[Routing-Configuration] Section | 340

[Supported-MAP-Messages] Section | 343

[Common-AGW-Configurations] Section | 344

[Process<name>] Section | 347

Configuring the authGateway Startup Information | 348

Example—Creating and Starting the authGateway Process | 352

Configuring the GWrelay.conf File | 353

Example—Configuring the GWrelay.conf File | 353

Starting the GWrelay Process | 353

Configuring the ulcmmg.conf File | 354

Example | 354

Loading the MML Configuration Settings | 354

Example | 355

Sample authGateway Command and File | 355

sampleCreateProcess.mml | 355

authGateway.conf (Gateway Routing Configuration File) | 355

## 7

## Uninstalling Steel-Belted Radius Carrier Software

Removing Steel-Belted Radius Carrier Software | 359

Uninstalling the Steel-Belted Radius Carrier Server Software | 359

## 8

## Upgrading Your SSR Cluster

Overview of Upgrading Your SSR Cluster | 366

Upgrade Methods | 366

Notes on Working with Data (D) Nodes | 367

Using the SSR Configuration Script | 368

## Using a Transition Server to Mitigate Downtime While Upgrading Your Cluster | 369

Cluster Migration Strategy | 369

Using a Transition Server | 370

Individual Node Migration Guidelines | 372

Cluster Migration Workflow | 372

Creating a Transition Server | 373

Preparing the Transition Server | 373

Unpacking and Configuring the New Software on the Transition Server | 374

Configuring the Transition Server | 385

Switching Traffic to the Transition Server | 387

Installing the Cluster | 387

Removing the Transition Server from Service | 390

Cleaning the Transition Server | 390

## Upgrading Your Cluster Using the Rolling Restart Method | 396

Overview of the Rolling Restart Upgrade Method | 396

Summary of the Rolling Restart Upgrade Method | 397

Introduction and Requirements | 397

Preparation | 399

Upgrading the M Nodes in the Cluster | 400

Installing the New SBRC Software on the M Nodes | 400

Configuring the New SBRC Software on the M Node | 404

Upgrading the SM Nodes in the Cluster | 409

Installing the New SBRC Software on the SM Node | 409

Configuring the New SBRC Software on the SM Node | 413

Upgrading the S Nodes in the Cluster | 422

Installing the New SBRC Software on the S Node | 422

Configuring the New SBRC Software on the S Node | 425

Upgrading the Data (D) Nodes with the New Software | 431

Installing the Complete SBR Release Distribution Package | 431

Installing Only the Latest MySQL Package | 435

Configuring the New Software on the Data Node | 439

Launching Web GUI | 446

## **Upgrading Your Cluster Using the Backup, Destroy, and Re-Create Method | 447**

Overview of the Backup, Destroy, and Re-Create Upgrade Method | 448

Capturing Your Current Cluster Configuration | 448

Backing Up the State of the Sessions in Your SSR Database | 450

Destroying the Cluster Database | 450

Stopping All Processes on All Nodes | 451

Installing the New SBRC Software on the Data Nodes | 452

- Uninstalling and Removing the Software Package on the Data Nodes | 452

- Installing the New Software Package on the Data Nodes | 454

Installing the New SBRC Software on the SM Nodes | 455

Configuring the First SM Node and Creating a New Cluster Definition | 457

Configuring the SBRC Software on the SM Nodes | 463

Configuring the SBRC Software on the Data Nodes | 469

Starting the SSR Processes | 473

Creating and Restoring the SSR Database | 476

Starting the RADIUS Process | 478

Launching Web GUI | 480

## **Part 9**

## **Appendix**

### **Appendix A**

## **Testing the Installation with DemoSetup.sh | 482**

Creating a Test Database | 482

DemoSetup.sh Syntax | 484

- Syntax | 484

- Options | 484

- Notes | 485

- Example | 485

# About This Guide

## IN THIS SECTION

- Objective | [xlvi](#)
- Audience | [xlvi](#)
- Documentation Conventions | [xlvii](#)
- Related Documentation | [I](#)
- Obtaining Documentation | [Ivi](#)
- Documentation Feedback | [Ivi](#)
- Requesting Technical Support | [Ivii](#)

This preface provides the following guidelines for using the *SBR Carrier Installation Guide*:

## Objective

This guide describes how to configure and administer Steel-Belted Radius Carrier software running on the Solaris operating system.

## Audience

This guide is intended for network administrators working for wireline and wireless carriers that are deploying converged services or emerging wireless technologies such as Worldwide Interoperability Microwave Access (WiMAX). It provides the information that administrators need to implement and maintain authentication, authorization, and accounting (AAA) services.

This guide assumes that you are familiar with general RADIUS and networking concepts, as well as the network environment that includes Steel-Belted Radius Carrier.

If you use Steel-Belted Radius Carrier with third-party products such as Oracle, this guide assumes you are familiar with the installation, configuration, and use of those products.

## Documentation Conventions

Table 1 on page xlviii defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
<b>NOTE:</b>	Informational note	Indicates important features or instructions.
<b>CAUTION:</b>	Caution	Indicates a situation that might result in loss of data or hardware damage.
<b>WARNING:</b>	Warning	Alerts you to the risk of personal injury.

Table 2 on page xlviii describes the text conventions used throughout this manual.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Text Conventions</b>		
<b>Bold text like this</b>	Represents commands and keywords in text.	<ul style="list-style-type: none"> <li>Issue the <b>clock source</b> command.</li> <li>Specify the keyword <b>exp-msg</b>.</li> </ul>
<b>Bold text like this</b>	Represents text that the user must type.	<b>host1(config)#traffic class low-loss1</b>
Fixed-width text like this	Represents information as displayed on your terminal's screen.	<pre>host1#show ip ospf 2  Routing Process OSPF 2 with Router ID 5.5.0.250  Router is an Area Bordèr Router (ABR)</pre>
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Emphasizes words.</li> <li>Identifies variables.</li> <li>Identifies chapter, appendix, and book names.</li> </ul>	<ul style="list-style-type: none"> <li>There are two levels of access, <i>user</i> and <i>privileged</i>.</li> <li><i>clusterId</i>, <i>ipAddress</i>.</li> <li><i>Appendix A, System Specifications</i>.</li> </ul>
Plus sign (+) linking key names	Indicates that you must press two or more keys simultaneously.	Press Ctrl+b.



Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
<i>radiusdir</i>	Represents the directory into which Steel-Belted Radius Carrier has been installed. The default location is <b>/opt/JNPRsbr/radius</b> on Solaris systems, but any location may be specified during installation.	Change directories to <b>/radiusdir/radiusdir</b>
<b>Syntax Conventions</b>		
Plain text like this	Represents keywords.	terminal length
<i>Italic text like this</i>	Represents variables.	<i>mask</i> , <i>accessListName</i>
< > (angle brackets)	Enclose a list of possible selections.	<add   replace>
(pipe symbol)	Represents a choice to select one keyword or variable in a list of choices that is separated by the pipe symbol.	diagnostic   line  In this example, you must specify <i>add</i> or <i>replace</i> but cannot specify both:  <add   replace>  Attribute [ ,Attribute]
[ ] (brackets)	Represent optional keywords or variables.	[ internal   external ], or <add   replace> = Attribute [,Attribute], where the second attribute is identified as optional by the brackets.  When they are used in a configuration files brackets identify a section of the file. In scripts or in operating system commands, brackets indicate the default response or entry.
[ ]* (brackets and asterisk)	Represent optional keywords or variables that can be entered more than once.	[ level1   level2   l1 ]*
{ } (braces)	Represent required keywords or variables.	{ permit   deny } { in   out } { clusterId   ipAddress }

## Related Documentation

Table 3 on page I lists and describes the Steel-Belted Radius Carrier documentation set:

Table 3: Steel-Belted Radius Carrier Documentation

Document	Description
<i>Steel-Belted Radius Carrier Installation Guide</i>	Describes how to install the Steel-Belted Radius Carrier software on the server.
<i>Steel-Belted Radius Carrier Administration and Configuration Guide</i>	Describes how to configure and operate the Steel-Belted Radius Carrier and its separately licensed modules.
<i>Steel-Belted Radius Carrier Reference Guide</i>	Describes the settings and valid values of the Steel-Belted Radius Carrier configuration files.
<i>Steel-Belted Radius Carrier Performance, Planning and Tuning Guide</i>	Provides tips, use cases, and tools you need to: <ul style="list-style-type: none"><li>• Improve SBRC performance through planning, analysis, and configuration</li><li>• Increase SBRC throughput and reliability</li><li>• Analyze specific use cases, in the lab or in the production environment, to identify areas of potential performance enhancement and to limit the impact of resource constraints and failure scenarios</li></ul>
<i>Steel-Belted Radius Carrier Release Notes</i>	Contains the latest information about features, changes, known problems, and resolved problems.

**NOTE:** If the information in the Release Notes differs from the information in any guide, follow the Release Notes.

### Requests for Comments (RFCs)

The Internet Engineering Task Force (IETF) maintains an online repository of Request for Comments (RFCs) online at <http://www.ietf.org/rfc.html>.

Table 4 on page li lists the RFCs that apply to Steel-Belted Radius Carrier.

Table 4: RFCs Related to Steel-Belted Radius Carrier

RFC Number	Title
RFC 1035	<i>Domain Names - Implementation and Specification</i> . P. Mockapetris. November 1987.
RFC 1155	<i>Structure and Identification of Management Information for TCP/IP-based Internets</i> . M. Rose, K. McCloghrie, May 1990.
RFC 1213	<i>Management Information Base for Network Management of TCP/IP-based internets: MIB-II</i> . K. McCloghrie, M. Rose, March 1991.
RFC 2006	<i>The Definitions of Managed Objects for IP Mobility Support using SMIv2</i> . D. Cong and others. October 1996.
RFC 2104	<i>HMAC: Keyed-Hashing for Message Authentication</i> . H. Krawczyk, M. Bellare, R. Canetti. February 1997.
RFC 2246	<i>The TLS Protocol</i> . T. Dierks, C. Allen. January 1999.
RFC 2271	<i>An Architecture for Describing SNMP Management Frameworks</i> . D. Harrington, R. Presuhn, B. Wijnen, January 1998.
RFC 2284	<i>PPP Extensible Authentication Protocol (EAP)</i> . L. Blunk, J. Vollbrecht, March 1998.
RFC 2433	<i>Microsoft PPP CHAP Extensions</i> . G. Zorn, S. Cobb, October 1998.
RFC 2548	<i>Microsoft Vendor-specific RADIUS Attributes</i> . G. Zorn. March 1999.
RFC 2607	<i>Proxy Chaining and Policy Implementation in Roaming</i> . B. Aboba, J. Vollbrecht, June 1999.
RFC 2618	<i>RADIUS Authentication Client MIB</i> . B. Aboba, G. Zorn. June 1999.
RFC 2619	<i>RADIUS Authentication Server MIB</i> . G. Zorn, B. Aboba. June 1999.
RFC 2620	<i>RADIUS Accounting Client MIB</i> . B. Aboba, G. Zorn. June 1999.
RFC 2621	<i>RADIUS Accounting Server MIB</i> . G. Zorn, B. Aboba. June 1999.
RFC 2622	<i>PPP EAP TLS Authentication Protocol</i> . B. Aboba, D. Simon, October 1999.

**Table 4: RFCs Related to Steel-Belted Radius Carrier (continued)**

RFC Number	Title
RFC 2719	<i>Framework Architecture for Signaling Transport</i> . L. Ong et al., October 1999
RFC 2809	<i>Implementation of L2TP Compulsory Tunneling via RADIUS</i> . B. Aboba, G. Zorn. April 2000.
RFC 2865	<i>Remote Authentication Dial In User Service (RADIUS)</i> . C. Rigney, S. Willens, A. Rubens, W. Simpson. June 2000.
RFC 2866	<i>RADIUS Accounting</i> . C. Rigney. June 2000.
RFC 2867	<i>RADIUS Accounting Modifications for Tunnel Protocol Support</i> . G. Zorn, B. Aboba, D. Mitton. June 2000.
RFC 2868	<i>RADIUS Attributes for Tunnel Protocol Support</i> . G. Zorn, D. Leifer, A. Rubens, J. Shriver, M. Holdrege, I. Goyret. June 2000.
RFC 2869	<i>RADIUS Extensions</i> . C. Rigney, W. Willats, P. Calhoun. June 2000.
RFC 2882	<i>Network Access Servers Requirements: Extended RADIUS Practices</i> . D. Mitton. July 2000.
RFC 2960	<i>Stream Control Transmission Protocol</i> . R. Stewart and others. October 2000.
RFC 3046	<i>DHCP Relay Agent Information Option</i> . M. Patrick. January 2001.
RFC 3118	<i>Authentication for DHCP Messages</i> . R.Droms and others. June 2001.
RFC 3162	<i>RADIUS and IPv6</i> . B. Aboba, G. Zorn, D. Mitton. August 2001.
RFC 3344	<i>IP Mobility Support for IPv4</i> . C. Perkins. August 2002.
RFC 3539	<i>Authentication, Authorization, and Accounting (AAA) Transport Profile</i> . B. Aboba, J. Wood. June 2003.
RFC 3575	<i>IANA Considerations for RADIUS (Remote Authentication Dial-In User Service)</i> . B. Aboba, July 2003.
RFC 3576	<i>RFC3576 - Dynamic Authorization Extensions to Remote to Remote Authentication Dial In User Service</i> . Network Working Group, 2003

**Table 4: RFCs Related to Steel-Belted Radius Carrier (continued)**

RFC Number	Title
RFC 3579	<i>RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)</i> . B. Aboba, P. Calhoun, September 2003.
RFC 3580	<i>IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines</i> . P. Congdon, B. Aboba, A. Smith, G. Zorn, J. Roese, September 2003.
RFC 3588	<i>Diameter Base Protocol</i> . P. Calhoun, J. Loughney, E. Guttman, G. Zorn, J. Arkko. September 2003.
RFC 3748	<i>Extensible Authentication Protocol</i> . B. Aboba, L. Blunk, J. Vollbrecht, J. Carlson, H. Levkowetz. June 2004.
RFC 3957	<i>Authentication, Authorization, and Accounting (AAA) Registration Keys for Mobile IPv4</i> . C. Perkins and P. Calhoun. March 2005.
RFC 4005	<i>Diameter Network Access Server Application</i> . P. Calhoun, G. Zorn, D. Spence, D. Mitton. August 2005.
RFC 4017	<i>Extensible Authentication Protocol (EAP) Method Requirements for Wireless LANs</i> . D. Stanley and others. March 2005.
RFC 4072	<i>Diameter Extensible Authentication Protocol (EAP) Application</i> . P. Eronen, G. Zorn, T. Hiller. August 2005.
RFC 4186	<i>Extensible Authentication Protocol Method for Global System for Mobile Communications (GSM) Subscriber Identity Modules (EAP-SIM)</i> . H. Haverinen, J. Salowey. January 2006.
RFC 4187	<i>Extensible Authentication Protocol Method for Global System for 3rd Generation Authentication and Key Agreement (EAP-AKA)</i> . J. Arkko, H. Haverinen. January 2006.
RFC 4282	<i>The Network Access Identifier</i> . B. Aboba and others. December 2005.
RFC 4284	<i>Identity Selection Hints for the Extensible Authentication Protocol (EAP)</i> . F. Adrangi, V. Lortz, F. Bari, P. Eronen. January 2006.
RFC 4306	<i>Internet Key Exchange (IKEv2) Protocol</i> . C. Kaufman. December 2005.

Table 4: RFCs Related to Steel-Belted Radius Carrier (continued)

RFC Number	Title
RFC 4372	<i>Chargeable User Identity</i> . F. Adrangi and others. January 2006.
RFC 4510	<i>Lightweight Directory Access Protocol (LDAP) Technical Specification Road Map</i> . K. Zeilenga, June 2006.
RFC 4666	<i>Signaling System 7 (SS7) Message Transfer Part 3 (MTP3) - User Adaptation Layer (M3UA)</i> . K. Morneault, J. Pastor-Balbas. September 2006.
RFC 4668	<i>RADIUS Authentication Client MIB for IPv6</i> . D. Nelson. August 2006.
RFC 4669	<i>RADIUS Authentication Server MIB for IPv6</i> . D. Nelson. August 2006.
RFC 4670	<i>RADIUS Accounting Client MIB for IPv6</i> . D. Nelson. August 2006.
RFC 4671	<i>RADIUS Accounting Server MIB for IPv6</i> . D. Nelson. August 2006.
RFC 5281	<i>Extensible Authentication Protocol Tunneled Transport Layer Security Authenticated Protocol Version 0 (EAP-TTLSv0)</i> . P. Funk, S. Blake-Wilson. August 2008.
RFC 5448	<i>Improved Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA')</i> . J. Arkko, V. Lehtovirta, P. Eronen. May 2009.
RFC 5997	<i>Use of Status-Server Packets in the Remote Authentication Dial In User Service (RADIUS) Protocol</i> . A. DeKok. August 2010.
RFC 6733	<i>Diameter Base Protocol</i> . V. Fajardo, J. Arkko, J. Loughney, G. Zorn. October 2012.
RFC 6911	<i>RADIUS Attributes for IPv6 Access Networks</i> . W. Dec, B. Sarikaya, G. Zorn, D. Miles, B. Lourdelet. April 2013.

### 3GPP Technical Specifications

The Third-Generation Partnership Project (3GPP) and 3GPP2 maintains an online repository of Technical Specifications and Technical Reports at <http://www.3gpp.org> and <http://www.3gpp2.org>, respectively.

Table 5 on page lv lists the 3GPP Technical Specifications that apply to Steel-Belted Radius Carrier.

Table 5: 3GPP Technical Specifications

3GPP TS Number	Title	Applicable Sections
3GPP TS 22.234 Version 12.0.0	<i>Requirements on 3GPP system to Wireless Local Area Network (WLAN) interworking</i>	<ul style="list-style-type: none"> <li>● Section 5.1.7: Interworking between PLMN and WLANs</li> </ul>
3GPP TS 23.003 Version 12.6.0	<i>Numbering, addressing, and identification</i>	<ul style="list-style-type: none"> <li>● Section 2.2: Composition of IMSI</li> </ul>
3GPP TS 23.008 Version 12.6.0	<i>Organization of subscriber data</i>	<ul style="list-style-type: none"> <li>● Section 3B: Definition of subscriber data I-WLAN domain</li> </ul>
3GPP TS 23.234 Version 12.0.0	<i>3GPP system to Wireless Local Area Network (WLAN) interworking; System description</i>	<ul style="list-style-type: none"> <li>● Section 6.1: Reference Model</li> <li>● Section 6.2: Network Elements</li> </ul>
3GPP TS 23.402 Version 12.8.0	<i>Architecture enhancements for non-3GPP accesses</i>	<ul style="list-style-type: none"> <li>● Section 4.1: Concepts</li> <li>● Section 4.3: Network Elements</li> </ul>
3GPP TS 24.302 Version 14.4.0	<i>Access to the 3GPP Evolved Packet Core (EPC) via non-3GPP access networks; Stage 3</i>	<ul style="list-style-type: none"> <li>● Section 6: UE – EPC Network protocols</li> <li>● Section 8: PDUs and parameters specific to the present document</li> </ul>
3GPP TS 29.002 Version 12.7.0	<i>Mobile Application Part (MAP) specification</i>	<ul style="list-style-type: none"> <li>● Section 6: Requirements concerning the use of SCCP and TC</li> <li>● Section 7.1: Terminology and definitions</li> <li>● Section 7.2: Modelling principles</li> <li>● Section 7.3: Common MAP service</li> </ul>
3GPP TS 29.273 Version 12.7.0	<i>Evolved Packet System (EPS); 3GPP EPS AAA interfaces</i>	<ul style="list-style-type: none"> <li>● Section 4: SWa Description</li> <li>● Section 5: STa Description</li> <li>● Section 6: SWd Description</li> <li>● Section 7: SWm Description</li> <li>● Section 8: SWx Description</li> <li>● Section 9: S6b and H2 Description</li> <li>● Section 10: Result-Code and Experimental-Result Values</li> </ul>

Table 5: 3GPP Technical Specifications (*continued*)

3GPP TS Number	Title	Applicable Sections
3GPP TS 33.402 Version 14.2.0	<i>3GPP System Architecture Evolution (SAE); Security aspects of non-3GPP accesses</i>	<ul style="list-style-type: none"> <li>• <i>Section 6: Authentication and key agreement procedures</i></li> <li>• <i>Section 7: Establishment of security contexts in the target access system</i></li> <li>• <i>Section 8: Establishment of security between UE and ePDG</i></li> <li>• <i>Section 9: Security for IP based mobility signalling</i></li> <li>• <i>Section 14: Temporary identity management</i></li> </ul>

### WiMAX Technical Specifications

The WiMAX Forum Networking Group (NWG) maintains a repository of technical documents and specifications online at <http://www.wimaxforum.org>. You can also view the WiMAX IEEE standards, 802.16e-2005 for mobile WiMAX and 802.16-2004 for fixed WiMAX, online at <http://www.ieee.org>.

### Third-Party Products

For information about configuring your Ulticom software and hardware, or your access servers and firewalls, consult the manufacturer's documentation.

## Obtaining Documentation

To obtain the most current version of all Juniper Networks technical documentation, see the products documentation page on the Juniper Networks website at <https://www.juniper.net/>.

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation to better meet your needs. Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net), or fill out the documentation feedback form at <https://www.juniper.net/documentation/feedback/>. If you are using e-mail, be sure to include the following information with your comments:

- Document name
- Document part number



- Page number
- Software release version

## Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- **JTAC Policies**—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>
- **Product Warranties**—For product warranty information, visit <https://www.juniper.net/support/warranty/>
- **JTAC Hours of Operation**—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

### Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings:  
<https://support.juniper.net/support/>
- Find product documentation:  
<https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes:  
<https://support.juniper.net/support/downloads/>
- Search technical bulletins for relevant hardware and software notifications:  
<https://kb.juniper.net/InfoCenter/index?page=subscriptions>, "Manage My Subscriptions"
- Open a case online in the Juniper Networks Customer Service Portal:  
<https://my.juniper.net>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool located at

<https://entitlementsearch.juniper.net/entitlementsearch/>

### Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Juniper Networks Customer Service Portal at <https://my.juniper.net>
- Call 1-888-314-JTAC (1-888-314-5822 – toll free in the USA, Canada, and Mexico)

For international or direct-dial options in countries without toll-free numbers, visit <https://www.juniper.net/support/requesting-support.html>

When you contact technical support, be ready to provide:

- Your Steel-Belted Radius Carrier release number (for example, Steel-Belted Radius Carrier Release 7.x).
- Information about the server configuration and operating system, including any OS patches that have been applied.
- For licensed products under a current maintenance agreement, your license or support contract number.
- A detailed description of the problem.
- Any documentation that may help in resolving the problem, such as error messages, memory dumps, compiler listings, and error logs.

# 1

PART

## Overview of Steel-Belted Radius Carrier and Session State Register Installation

---

[Standalone Steel-Belted Radius Carrier Installation Overview](#) | 2

[Session State Register Cluster Installation Overview](#) | 9

---

# Standalone Steel-Belted Radius Carrier Installation Overview

## IN THIS CHAPTER

- [Migrating from Earlier SBR Releases | 2](#)
- [Supporting SIM with Signalware | 3](#)
- [Installation Workflow | 4](#)
- [Release 8.6.0 Standalone Server Installation Checklist | 4](#)

This chapter presents an overview of the installation process for standalone Steel-Belted Radius Carrier (SBRC) servers that do not use the Session State Register option and are not members of a Session State Register (SSR) cluster. The chapter includes a checklist of required tasks that you can use to track your progress as you work through the installation process.

For introductory information about installing a Session State Register cluster, see [“Session State Register Cluster Installation Overview” on page 9](#).

This chapter includes these topics:

## Migrating from Earlier SBR Releases

Migration involves moving a copy of the old SBR installation into a directory on the new Release 8.6.0 server. When Steel-Belted Radius Carrier is installed, files that may be moved forward are copied and integrated into the new installation by the configuration script. The corresponding Release 8.6.0 files are loaded on the server too, but not activated.

**NOTE:** Migration to SBR Release 8.6.0 is not supported from certain SBR versions earlier than Release 7.2.0.

**NOTE:** Migration is supported between like operating systems only—one Oracle or Solaris server to a second and from the same type of server CCM role (primary or replica) to the same type of server. Cross-OS and cross-type migrations are not possible.

Many older release configuration files may not be moved forward because they conflict with newer versions. However, portions of these files may be copied and pasted into the corresponding Release 8.6.0 files to preserve existing settings without disabling new features.

**NOTE:** In SBR Carrier Release 7.0, you configured the **ss7ldapdb.gen** file to store CDR records in the LDAP directory. When you run SBR Carrier Release 8.6.0 in standalone mode, these settings are configured in the **ss7db.gen** file. When you migrate from Release 7.0 to Release 8.6.0, you need to move these settings from the **ss7ldapdb.gen** file to the **ss7db.gen** file.

Previous releases of SBRC software running in standalone mode and using the SIM module used a different method to store the CDR records; SBRC Release 8.6.0 uses an LDAP directory to store these records. This directory is configured with the **ss7db.gen** file. After you upgrade to SBRC Release 8.6.0, the CDR records from the previous releases of SBR Carrier software are no longer accessible.

For details about what works when moving from one of the different supported releases, see [“Migrating from Previous SBR Releases” on page 56](#).

## Supporting SIM with Signalware

Steel-Belted Radius Carrier support for SIM and other RADIUS protocols that use EAP-SIM or EAP-AKA (such as WiMAX, which may use EAP-AKA), depend on Signalware 9 with Service Pack 6.A software being properly installed on the server.

**NOTE:** We strongly recommend that Signalware software be installed and configured before Steel-Belted Radius Carrier is installed because the SBR Carrier installation incorporates Signalware settings, if they are available.

Juniper Networks provides an example Signalware installation with:

- The correct settings for simple integration with Steel-Belted Radius Carrier.
- The recommended procedures to configure Signalware 9 to support Steel-Belted Radius Carrier.

- The recommended procedure for migrating configuration settings forward from an existing Signalware installation.

See [“SIGTRAN Support for Steel-Belted Radius Carrier” on page 309](#).

**NOTE:** Because the Signalware communication stack is not supported on Solaris, the SIM authentication module cannot be used on the Solaris platform to communicate with an HLR to process RADIUS requests. However, the module can be used with an HSS by using the RADIUS to Diameter conversion feature.

## Installation Workflow

The installation workflow follows a list of tasks that must be performed to make the Steel-Belted Radius Carrier server fully functional. This guide presents groups of related tasks in sequential order.

The high-level view of the installation has three groups of tasks:

1. **Before You Begin**—The tasks to complete and the information to gather before installing the Steel-Belted Radius Carrier software. These topics are discussed in detail in [“Before You Install Software” on page 39](#), and in the succeeding chapters that document the optional installation of Signalware that is required to support some features.
2. **Software installation**—Unpacking the software package, selecting configuration tasks, and providing settings that are used to set up the server. These topics are discussed in detail in [“Installation and Basic Configuration of a SBR Carrier Standalone Server” on page 64](#).
3. **Basic configuration**—Starting and stopping the server daemon, setting up SNMP, and editing configuration files to include information from earlier releases. These topics are discussed in detail in [“Basic SBR Carrier Configuration” on page 84](#).

## Release 8.6.0 Standalone Server Installation Checklist

This section includes an installation checklist for a standalone Steel-Belted Radius Carrier (not using Session State Register). Use the checklist to record the steps for a new installation. It includes all possible options, including integrating configuration files from a previous SBR release. If your installation does not use all the available options, you can skip those steps.

This set of procedures installs a new, unconfigured Steel-Belted Radius Carrier Release 8.6.0 build on a Solaris platform. It includes all possible options that can be included in the Steel-Belted Radius Carrier installation.

## Before You Begin

Before you begin to install the Steel-Belted Radius Carrier software, perform these preliminary tasks or verify that they have been accomplished:

1. Review the Steel-Belted Radius Carrier *Release Notes* for up-to-date information that became available after the Steel-Belted Radius Carrier guides were published.

Download the most recent Release Notes file from:

[https://www.juniper.net/documentation/en\\_US/sbr-carrier8.6.0/information-products/topic-collections/release-notes/sbrc-8.6.0.pdf](https://www.juniper.net/documentation/en_US/sbr-carrier8.6.0/information-products/topic-collections/release-notes/sbrc-8.6.0.pdf)

If the information in the *Release Notes* differs from the information in any Steel-Belted Radius Carrier guide, follow the *Release Notes*.

2. Ensure that the server chassis is physically secure.

To review recommendations on the server's characteristics, see ["Selecting an Appropriate Server" on page 41](#).

## Migration

3. If you are migrating from an earlier SBR server release, put a copy of that installation on the system that you want to host Release 8.6.0.

The recommended method is:

- a. As root on the old server, shut down the SBR daemon.

Execute: `/radiusdir /sbrd stop`.

- b. Create a gzip archive of the *radiusdir* directory, preserving the file structure.

Change directories to `/opt/JNPRsbr/` and execute (for example):

**`tar cfE - radius |gzip > 6_radius.tgz`**

- c. As root on the Release 8.6.0 host, create a working directory for the archive:

Execute:

**`mkdir -p /opt/JNPRsbr/PreviousRelease_backup`**

and:

**`cd /opt/JNPRsbr/PreviousRelease_backup`**

- d. Use FTP binary mode to copy the archive file to the Release 8.6.0 host.
- e. Gunzip the archive in the working directory. This does not create a working installation, but serves as a source of files that can be reused by the Release 8.6.0 server.

Execute (using the example file name):

```
gunzip 6_radius.tgz |tar xf -
```

- f. Check the file permissions on the unzipped archive to ensure that the files are writable.

See [“Creating a Copy of Existing SBR Server Release Files for Migration” on page 54](#) for more information.

4. Confirm that the server chassis provides at least the minimum hardware and software requirements.

See [“Meeting System Requirements” on page 41](#).

5. Verify root access.

6. Verify that the server meets basic network requirements:

- See [“Verifying Network Connectivity” on page 49](#).
- See [“Verifying Hostname Resolution” on page 50](#).

## CCM

7. Verify the server role in the centralized configuration management (CCM) environment. The installation script can configure a Steel-Belted Radius Carrier server as:

- An autonomous server.
- The primary server supplying a group of replication servers.
- A replication server.

In an environment that uses replication servers, the primary server must be configured first to provide a base for replication.

## LDAP

8. (Optional) If the server will interact with a Lightweight Directory Access Protocol (LDAP) directory service agent, verify or record the path to the LDAP library files. (The default path is **/usr/lib**.)
9. (Optional) If the Steel-Belted Radius Carrier server will authenticate through an external database, you need to:

- Verify database compatibility and interoperability.

See [“Setting Up External Database Connectivity \(Optional\)” on page 53](#).

## ORACLE

10. If the external database is Oracle, the Steel-Belted Radius Carrier server must be configured as a client of the Oracle server.

As you set up the client, collect this information, or if the client is already installed, record this information to use while installing SBR software:



- Record the path to the local Oracle Home directory (for example:  
`/opt/10g/app/oracle/product/10.2.0.3`).
- Record the path to the local Oracle shared library (for example:  
`/opt/10g/app/oracle/product/10.2.0.3/lib`)
- Record the path to the local TNS\_ADMIN (for example:  
`/opt/10g/app/oracle/product/10.2.0.3/network/admin`)

## SIM / WiMAX / SIGTRAN

12. (Required only for these modules) Install and configure Signalware 9 with Service Pack 6.A software, following the product documentation and the guidelines in [“SIGTRAN Support for Steel-Belted Radius Carrier” on page 309](#). This step is mandatory for the SIM authentication module and any other SIGTRAN applications. This step is only required for the optional WiMAX mobility module if it uses the EAP-AKA authentication protocol, which requires the optional SIM authentication module.

## Installing SBR Software

If you have completed all items on the “Before You Begin” list, you are ready to install the SBR software. These are the key steps:

1. Unpack the Steel-Belted Radius Carrier Software.

See [“Unpacking the Steel-Belted Radius Carrier Software” on page 64](#).

If you are not familiar with UNIX package management commands, you can review a list of [“Package Management Commands” on page 64](#).

2. Check for running SNMP processes on the server and shut down any that are active.
3. Run the Steel-Belted Radius Carrier **configure** script. As the script runs, enter the information you recorded on the previous pages.

See [“Setting Up a Starter Kit’s First SBR/Management Node” on page 99](#).

## Basic Configuration

After the Steel-Belted Radius Carrier software is on the Solaris server, ensure that the base software runs properly and then perform basic configuration steps:

1. Start the Steel-Belted Radius Carrier daemon.
  - For standalone systems that are not in a Session State Register cluster:

The basic commands to start and stop the daemon from the command line are as follows:

- Start the RADIUS Server:  
`/radiusdir/sbrd start`

or restart: `/radiusdir/sbrd restart`

- Stop the RADIUS Server:  
`/radiusdir/sbrd stop`
- Display RADIUS Status Information:  
`/radiusdir/sbrd status`

For more information, see [“Starting and Stopping a Standalone Steel-Belted Radius Carrier Server” on page 77](#).

- For systems that are part of a Session State Register cluster, starting SBR nodes is part of starting and stopping the entire cluster.

## 2. Launch the Web GUI.

Use an HTML browser to address the Steel-Belted Radius Carrier server HTML server. Download the Administrator client application and log in as root. Also confirm that the links to the documentation work.

Verify that you can download and run the application both from the local host and from a workstation on the network.

## Migrating from a Previous Release

### 3. Configure the environment.

If you installed a copy of a previous SBR server installation on the new server as mentioned in [“Migration” on page 5](#), the configuration files from that release are used in this startup. Be aware that to use some new features delivered in Release 8.6.0 you must edit these configuration files to include new settings. See [“Migrating from Previous SBR Releases” on page 56](#).

For a new installation, configuring the environment includes basic low-level configuration tasks such as:

- Configuring communications between any additional devices and the Steel-Belted Radius Carrier server.
- Adding additional users to the server.

See [“Configuring the Server” on page 86](#).

### 4. Configure SNMP. If you elected not to configure the Steel-Belted Radius Carrier SNMP agent during software installation, but have changed your mind and other agents also run on the server (such as the Solaris operating system agent), you must adjust the ports the agents use so multiple agents do not contend for the same resource.

See [“Configuring SNMP” on page 86](#).

# Session State Register Cluster Installation Overview

## IN THIS CHAPTER

- [SSR Cluster Overview | 9](#)
- [SSR Cluster Concepts and Terminology | 10](#)
- [Supported SBR Carrier SSR Cluster Configurations | 16](#)
- [Session State Register Database Tables | 22](#)
- [Supporting SIM, SIGTRAN, and Other Protocols with Signalware | 25](#)
- [Cluster Installation Workflow | 25](#)
- [Session State Register Cluster Installation Checklist | 26](#)

This chapter contains an overview of the Session State Register (SSR) and of the installation process for all nodes and servers in SSR cluster. It includes a checklist of required tasks that you can use to track your progress as you work through the installation process.

For introductory information about installing a standalone Steel-Belted Radius (SBR) Carrier server, see [“Standalone Steel-Belted Radius Carrier Installation Overview” on page 2](#).

This chapter includes these topics:

## SSR Cluster Overview

A standalone SBR Carrier server runs the entire AAA process using local session database processes on a single machine.

The Session State Register (SSR) module implements a stateless, highly reliable and highly available AAA platform. It separates SBR Carrier front end processes from back end data functions that take place on two or four Session State Register data servers.

The SBR Carrier front end and SSR back end servers collaborate to create a virtual AAA server that provides:

- High availability
- Session state preserved during failover of front end SBR Carrier nodes

- Application session awareness
- Centralized IP address management
- Concurrency management (if the optional Concurrency and Wholesale Module is installed on all SBR Carrier servers)

To work efficiently and to provide redundancy, a production SSR cluster must be built on a fast, isolated, and redundant network infrastructure. It must include multiple hosts and nodes so a single node failure cannot prevent the cluster from operating.

Because the SSR database contains detailed attributes for each session it serves, it is a logical source for applications that query the AAA environment for correlative and verification purposes. The default SSR database schema addresses standard requirements for most wireless and wireline service providers. You can also customize it to address unique requirements or to support other applications. For example, a video portal or streaming server that has the IP address for a session but requires the user identity attached to the session for billing or personalization. The video gateway can query the SSR database in real time to correlate information about any active session.

**NOTE:** The Session State Register Database is not, and cannot be used as, a permanent subscriber database.

RADIUS Class Attributes are part of the RADIUS standard, but not all network access devices (NADs) fully support them. In order for the SSR to reliably track sessions throughout each session's lifetime, the network NADs must support the RADIUS Class Attributes.

## SSR Cluster Concepts and Terminology

A Session State Register cluster has both a physical and logical organization. The physical elements are *servers*. The logical elements are *nodes*. The two terms are not interchangeable.

### Session State Register Servers

The Session State Register has requirements for the entire cluster and all servers that participate in the cluster, over and above the requirements for standalone SBR Carrier servers. So a SSR cluster does not have any single point of failure, each server in a cluster must have its own memory and disks. We do not recommend or support virtual servers, network shares, network file systems, and SANs.

All servers in the cluster require at least two physical Ethernet ports that provide the same throughput. Multipathing the NICs to a single IP address is required. Session State Register Cluster can work over a 100Base-T network but we recommend 1000Base-T (gigabit Ethernet).

All data servers must have equal processor power, memory space, and available bandwidth because they are tightly coupled and share data. If the overall throughput of the data servers varies from machine to machine, performance degrades. SBR Carrier servers and management servers' configuration may vary from machine to machine, so long as the basic standalone requirements are met.

## Session State Register Nodes

Four types of nodes can be included in a cluster, each with a specific role within the cluster:

- An *SBR node*, also known as an (*s*) node, hosts the SBR Carrier RADIUS process software component, any optional modules, and all related processes that read and write data into the SSR database. This type of node accesses and manipulates the cluster's shared data that is hosted by the data nodes.
- An *SBR/management node*, also known as an (*sm*) node, hosts a combination of a *SBR node* and a *management node*. The *SBR node* software component runs the SBR Carrier RADIUS process and the *management node* software component runs the management node SSR process.
- A *management node*, also known as an (*m*) node, controls itself and all data nodes in the cluster. It provides configuration data, starts and stops nodes, can back up the database and perform other database operations. It also manages a database process that supports the SSR storage engine. Cluster configuration data is located in an identical **config.ini** file on each of the cluster's management nodes.
- A *data node*, also known as a (*d*) node, runs the **ndbmtdd** data process. The **ndbmtdd** process cooperatively manages, replicates, and stores data in the SSR storage engine with other data nodes. Each data node has its own memory and permanent storage. Each one maintains both a portion of the working copy of the SSR database and a portion of one or more replicas of the database.

All the data nodes in a cluster run a special process called the *shared memory engine* that manages the working copy of the SSR database. The management nodes coordinate the service among the participating data nodes. The shared storage engine and the SSR database replaces the on-board database used by standalone Steel-Belted Radius Carrier servers. The shared memory engine ensures that the database is updated by a synchronous replication mechanism that keeps cluster nodes synchronized: a transaction is not committed until all cluster nodes are updated.

## SSR Data Entities

Each data node participates in a *node group* of two data nodes. A Starter Kit cluster has a single node group with two members; a Starter Kit with an Expansion Kit has two node groups, each with two data nodes. Each node group stores different partitions and replicas.

- A *partition* is a portion of all the data stored by the cluster. There are as many cluster partitions as node groups in the cluster. Each node group keeps at least one copy of any partitions assigned to it (that is, at least one replica) available to the cluster.

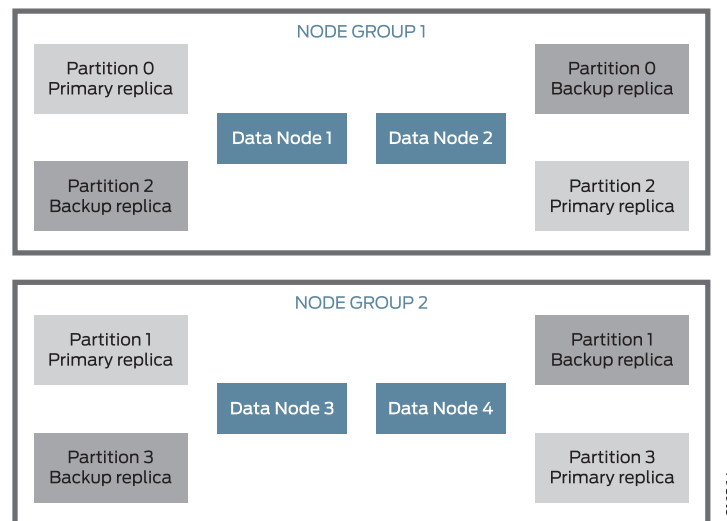
- A *replica* is a copy of a partition. Each data node in a node group stores a replica of a partition. A replica belongs entirely to a single data node; a node can (and usually does) store several replicas because maintaining two replicas is the fixed setting for SSR.

Figure 1 on page 12 shows the data components of a data cluster with four data nodes arranged in two node groups of two nodes each. Nodes 1 and 2 belong to Node Group 1. Nodes 3 and 4 belong to Node Group 2.

- Because there are four data nodes, there are four partitions.
- The number of replicas is two, to create two copies of each primary partition.

So long as both nodes in one node group are operating, or one node in each node group is operating, the cluster remains viable.

Figure 1: SSR with Four Data Nodes in Two Groups



The data stored by the cluster in Figure 1 on page 12 is divided into four partitions: 0, 1, 2, and 3. Multiple copies of each partition are stored within the same node group. Partitions are stored on alternate node groups:

- Partition 0 is stored on Node Group 1. A primary replica is stored on Data Node 1 and a backup replica is stored on Data Node 2.
- Partition 1 is stored on the other node group, Node Group 2. The primary replica is on Data Node 3 and its backup replica is on Data Node 4.
- Partition 2 is stored on Node Group 1. The placement of its two replicas is reversed from that of Partition 0; the primary replica is stored on Data Node 2 and the backup on Data Node 1.
- Partition 3 is stored on Node Group 2, and the placement of its two replicas are reversed from those of partition 1: the primary replica is on Data Node 4 and the backup on Data Node 3.

**TIP:** *Primary* and *replica* are used in another context in the Steel-Belted Radius Carrier environment and documentation, which can cause some confusion. These terms mean something specific in the context of Session State Register, but they are also used when talking about centralized configuration management, or CCM.

CCM is a feature that coordinates Steel-Belted Radius Carrier server settings between a primary RADIUS server and one or more replica RADIUS servers. It copies critical configuration files from the primary to the replicas, so it keeps multiple SBR Carrier servers operating the same way.


CCM is a separate tool and process that are not tied or linked to SSR, but CCM is often used in SSR environments to keep the SBR Carrier nodes operating identically.

## Cluster Configurations

For the highest level of redundancy, we recommend that each node in a cluster run on its own server. In many locations and for many installations, that might not be practical, so you can run a SBR and a management node together on the same machine—in fact, that is the default configuration for the SSR Starter Kit cluster. However, neither a management node nor an SBR node can run on the same machine as a data node. Separation is required so that management arbitration services continue if one of the data node servers fails.

Using these separation guidelines, the recommended minimum size of a Session State Register cluster is four physical machines: two machines that each run a SBR/management node combination, and two machines to host the data nodes. This configuration supports all licenses and nodes included in the Session State Register Cluster Starter Kit and is shown in [Figure 2 on page 13](#):

**Figure 2: Basic Session State Register Starter Kit Cluster**



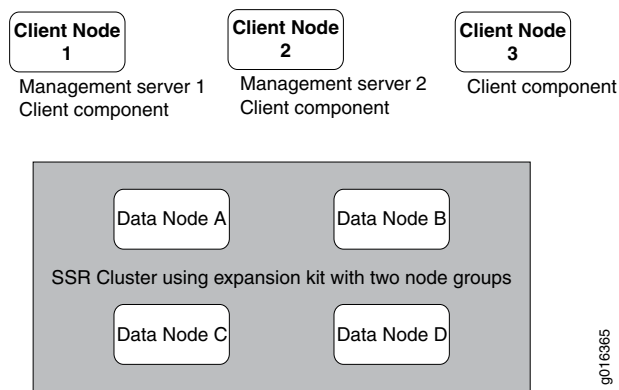
## Session State Register Scaling

You scale a Session State Register cluster when you add a separately licensed SSR Expansion Kit to a Starter Kit, a third management node, or additional SBR Carrier front end systems.

### ***Adding a Data Node Expansion Kit***

An Expansion Kit adds two data nodes to increase the number of data nodes in a cluster to four. The additional nodes form a second node group (as shown in [Figure 3 on page 14](#)) that provides more working memory for the SSR shared database. With the Expansion Kit in place, each node group manages a partition of the primary database and replicas. The data in each partition is synchronously replicated between the group's data nodes, so if one data node fails, the remaining node can still access all the data. This configuration also provides very quick failover times if a node fails.

**Figure 3: SSR Cluster with an Expansion Kit Setup to Create Two-Node Groups**



### ***Adding a Third Management Node***

A Management Node Expansion Kit provides software and a license for a third management node. If it is set up on a separate host instead of alongside a SBR Carrier node on a shared server, this also increases the resiliency of the cluster by providing an additional arbiter in case of a node failure.

### ***Adding More SBR Carrier Front End Servers***

The service capacity of the SBR Carrier environment grows when you add additional stateless SBR servers to the front end. Adding additional SBR Carrier servers increases the resiliency of the cluster and the speed of processing a particular transaction because wait time is reduced. Up to 20 Steel-Belted Radius Carrier nodes can be supported by a data cluster.

The SBR Carrier servers do not require identical configurations; they can be configured with different optional modules or communications interfaces. Each one requires a separate SBR Carrier license, but they all share the Session State Register Starter Kit license.

We recommend installing a load balancer in front of the SBR Carrier servers to evenly distribute the RADIUS load between front end SBR Carrier nodes. Regular server-based load balancing works if the front ends only processes RADIUS transactions, Use a RADIUS-aware load balancer if the front ends perform multi-round authentication.

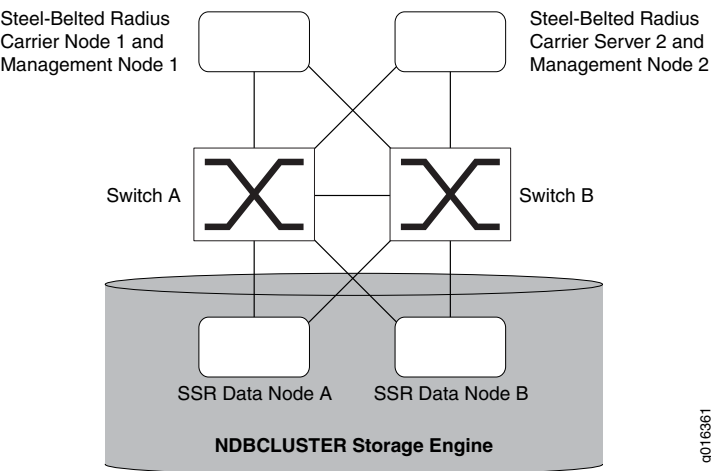


Cluster Network Requirements

A redundant cluster requires a redundant network. At the machine level, we require dual interface cards in each machine and multipathing.

We recommend that the network be a dedicated subnet with dual switches. This fully duplicates the network and each machine in the cluster has at least two routes to all other machines, as shown in [Figure 4 on page 15](#).

Figure 4: Starter Kit SSR Cluster with Redundant Network



The SSR database schema uses primary key lookups as often as possible during transaction processing, so the database cluster performance scales almost linearly based on the number of data nodes in the cluster.

Do not configure the subnet to be shared beyond the cluster machines because communications between nodes are not encrypted or shielded in any way. The only means of protecting transmissions within a cluster is to run your cluster on a protected network; do not interpose firewalls between any of the nodes.

Running the cluster on a private or protected network also increases efficiency because the cluster has exclusive use of all bandwidth between cluster hosts. This protects the cluster nodes from interference caused by transmissions between other devices on the network.

We strongly recommend Gigabit Ethernet as the network type; 100Base-T is the minimum supported speed. Network latency can severely degrade performance, so we also recommend that all servers be close enough together that latency is always much less than 10 ms.

Table 6: Latency between Servers and Its Effect on Performance

Latency Times	Performance Degradation
0 ms latency (LAN)	Baseline performance as designed.
10 ms latency	Up to 40% performance loss

Table 6: Latency between Servers and Its Effect on Performance (*continued*)

Latency Times	Performance Degradation
20 ms latency	Up to 60% performance loss
More than 20 ms latency	Not supported.

## Supported SBR Carrier SSR Cluster Configurations

A Starter Kit licenses you for two management node processes and two data node processes. SBR node processes are not included with the Starter Kit and must be purchased separately. You can collocate one or both SBR node processes with the management node processes, thereby creating one or two SBR/management combination nodes. You can add a third management node process to the cluster using the Management Node Expansion Kit. You always add data nodes in pairs. One Data Node Expansion Kit includes two data nodes.

**NOTE:** Each node type (s), (sm), (m), (d) requires its own machine.

You can have no more than three management node processes in a cluster.

The minimum requirements for a cluster are two SBR node processes (full licenses) and a Starter Kit license. You can deploy these licenses in a configuration of four, five, or six machines, depending on how many SBR node processes are collocated with management node processes on a single machine. Each SBR node process is licensed separately; you cannot share SBR licenses. Starter Kit licenses are shared by all processes in the cluster (through the **configure** script). [Table 7 on page 17](#) lists the possible configurations for the minimum requirements of two SBR node licenses and a Starter Kit license.

If all add-on products are added to the Starter Kit cluster, the maximum size of a data cluster is four data nodes, three management nodes, and up to 20 SBR Carrier nodes, as shown in [Table 7 on page 17](#).



**CAUTION:** Setting up an unsupported configuration can put data and equipment at risk and is not supported by Juniper Networks.

Also, note the latency limitation in [Table 6 on page 15](#). We do not support cluster configurations with latency between nodes that exceeds 20 ms, as can occur if servers are set up to spread a cluster across widely separated locations.

Table 7: Supported Cluster Configurations

Licenses	Node Type (One Machine Required for Each Node Type)			
	(S)	(SM)	(M)	(D)
Configuration 1:  Two SBR node processes (full licenses) and one Starter Kit (minimum configuration)	-	Two	-	Two
Configuration 2:  Two SBR node processes (full licenses) and one Starter Kit	One	One	One	Two
Configuration 3:  Two SBR node processes (full licenses) and one Starter Kit	Two	None	Two	Two
Configuration 1, 2, 3 with one Data Expansion Kit	-	-	-	Up to Four
Configuration 1, 2, 3 with one Management Node Expansion Kit	-	-	Up to Three	-
Maximum configuration:  Any of the previously listed configurations and additional SBR Nodes (front ends)	Up to a total of 20	-	Up to Three	Up to Four

## Failover Overview

To continue functioning without a service interruption after a component failure, a cluster requires at least 50 percent of its data and management nodes to be functional. If more than 50 percent of the data nodes fail, expect a service interruption, but continued operation of the available nodes.

Because SBR Carrier nodes function as front ends to the data cluster, they are not involved in any failover operations performed by the data cluster. However, as an administrator, you need to ensure that the front end environment is configured so that it can survive the loss of SBR Carrier nodes. (We recommend using an emergency IP address pool and running with a RADIUS-aware load balancer.)

A data cluster prepares for failover automatically when the cluster starts. During startup, two events occur:

- One of the data nodes (usually the node with the lowest node ID) becomes the *primary* of the node group. The primary node stores the authoritative copy of the database.
- One data node or management node is elected *arbitrator*. The arbitrator is responsible for conducting elections among the survivors to determine roles in case of node failures.

In a cluster, each management node and data node is allocated a vote that is used during this startup election and during failover operations. One management node is selected as the initial arbitrator of failover problems and of elections that result from them.


Within the cluster, data and management nodes monitor each other to detect communications loss and heartbeat failure. When either type of failure is detected, as long as nodes with more than 50 percent of the votes are operating, there is instantaneous failover and no service interruption. If exactly 50 percent of nodes and votes are lost, and if a data node is one of the lost nodes, the cluster determines which half of the database is to remain in operation. The half with the arbitrator (which usually includes the primary node) stays up and the other half shuts down to prevent each node or node group from updating information independently and then restarts and attempts to rejoin the active half of the cluster.

When a failed data node (or nodes) returns to service, the working nodes resynchronize the current data with the restored nodes so all data nodes are up to date. How quickly this takes place depends on the current load on the cluster, the length of time the nodes were offline, and other factors.

### Failover Examples

The following examples are based on the basic Starter Kit deployment setup with the recommended redundant network as shown in [Figure 5 on page 18](#). The cluster is set up in a single data center on a fully switched, redundant, layer 2 network. Each of the nodes is connected to two switches using the Solaris IP-multipathing feature for interface failover. The switches have a back-to-back connection.

**Figure 5: Starter Kit SSR Cluster with Redundant Network**



### **Possible Failure Scenarios**

With these basic configurations, a high level of redundancy is supported. So long as one data node is available to one SBR Carrier node, the cluster is viable and functional.

- If either SBR Carrier Server 1 or 2, which also run the cluster's management nodes, (s1 and m1 or s2 and m2) goes down, the effect on the facility and cluster is:
  - No AAA service impact.
  - NADs (depending on the failover mechanism in the device) switch to their secondary targets—the remaining SBR Carrier Server. Recovery of the NAD when the SBR Carrier Server returns to service depends on NAD implementation.
- If either data node A or B goes down, the effect is:
  - No AAA service impact; both SBR Carrier nodes continue operation using the surviving data node.
  - The management nodes and surviving data node detect that one data node has gone down, but no action is required because failover is automatic.
  - When the data node returns to service, it synchronizes its NDB data with the surviving node and resumes operation.
- If both management nodes (m1 and m2) go down, the effect is:
  - No AAA service impact because the all s and d nodes are still available. The data nodes continue to update themselves.
- If both data nodes go down, the effect on:
  - The management nodes is minimal. They detect that the data nodes are offline, but can only monitor them.
  - The SBR Carrier nodes varies:
    - Authentication of users and accounting for users that do not require shared resources such as the IP address pool or concurrency continue uninterrupted. If the nodes have local non-shared emergency IP address pools, the front ends can continue to process some requests.
    - Users that require shared resources are rejected.

The carrier nodes continue to operate this way until the data cluster comes back online; the cluster resumes normal AAA operation using the data cluster automatically.
- If one half of the cluster (SBR Carrier Server 1, management node 1, and data node A or SBR Carrier Server 2, management node 2, and data node B) goes down, the effect is:
  - No AAA service impact because the SBR Carrier node, a management node, and a data node are all still in service. NADs using the failed SBR Carrier Server fail over SBR Carrier Server.
  - When the failed data node returns to service, it synchronizes and updates its NDB data with the surviving node and resumes operation.

- When the failed SBR Carrier Server returns to service, the NADs assigned to use it as a primary resource return to service depending on the NAD implementation.

### ***Distributed Cluster Failure and Recovery***

You can divide a cluster and separate two equal halves between two data centers. In this case, the interconnection is made by dedicated communications links (shown as bold lines in [Figure 6 on page 21](#) and [Figure 7 on page 21](#)) that may be either:


- A switched layer 2 network, just as the single site cluster is set up.
- A routed layer 3 network that uses a routing table with backup routes to route over multiple links between data centers

However, this creates a configuration that is vulnerable to a catastrophic failure that severs the two halves of a dispersed cluster. We recommend adding a third management node at a location that has a separate alternative communication route to each half. A third management node:

- Eliminates the possibility of the cluster being evenly split by a communications failure.
- Creates an odd number of votes for elections, which greatly reduces the need for arbitration.


With a third management node in place, failover in the dispersed cluster is well managed because one side of the cluster does not have to determine what role to assume. Recovery is likely to be quicker when the data nodes are reunited because each node's status is more likely to have been monitored by at least one management node that is in communication with each segment.

Figure 6: Starter Kit SSR Cluster Divided Between Two Sites with Tertiary Management Node



Without a third management node, the configuration shown in [Figure 7 on page 21](#) is vulnerable to data loss if both communication links are severed or if the nodes in the primary half of the cluster all go offline simultaneously.

Figure 7: Starter Kit SSR Cluster Evenly Divided Between Two Sites



If either of those calamities occur, exactly half the nodes in a side survive. If the primary nodes are operating on one or both sides, the cluster continues to function. But the secondary side cannot determine whether

the primary side is really no longer available because it only has two votes. It can take 10-15 minutes for the secondary side of the cluster to automatically restart, promote itself to primary status, and resume cluster operations.

However, the SBR Carrier nodes that use that portion of the cluster do not automatically reconnect, and cannot communicate with the other half. If they were to reconnect to the secondary side, modifications made to the database create a divergence between the two copies of the database (although SBR Carrier nodes continue to process requests that do not require the database). The longer the cluster is split, the greater the divergence, and the longer it takes to resolve when recovery takes place.

To eliminate these problems, we recommend a proven alternative: adding a third management node in a third location that can communicate with each half of the dispersed cluster. Without the tertiary management node, there is a possibility of downtime in a dispersed cluster that suffers a catastrophic failure.

If you cannot add a third management node, we recommend that you configure the secondary side of the cluster not to automatically restart, but to go out of service when it instantaneously disconnects from the primary side nodes. Then you can determine the best course of action—to keep the cluster offline or to promote the secondary side of the cluster, relink the SBR Carrier nodes, and be aware that reconciling the divergence must be part of the recovery procedure.

When the cluster is reunited and goes into recovery mode, the primary and backup data nodes attempt to reconcile the divergence that occurred during separation. The moment they come in contact, transitory failures appear on the SBR Carrier nodes because the cluster configuration has changed; any transactions that are pending at that moment are terminated. The SBR Carrier nodes retry those transactions because they are classified as temporary failures; in most situations they are accepted on the first retry.

## Session State Register Database Tables

The database cluster uses a relational SQL design to store data. The database contains multiple tables that store Steel-Belted Radius Carrier state information such as IP address pool information and current sessions.

### IP Address Pools

Session State Register stores information about all IP address pools centrally in a dedicated table in the SSR database. Network access devices can send RADIUS requests to any stateless SBR Carrier server, and the SBR Carrier server allocates IP addresses from the shared central pool in response, after authentication.

Operators can query from any SSR management node to verify what addresses are in use and how many addresses remain available for an IP address pool or set of pools.



## Subscriber Session Data Controls

Session State Register stores session data centrally in the Current Sessions Table in the SSR database.

The table's format is controlled by a configuration file on the cluster's management nodes; one file is copied to all nodes in a node group, so all nodes operate with the same information. The standard Session State Register database schema addresses the needs of most carriers, but you can modify the CST to address unique needs and situations. The **CurrentSessions.sql** file determines the schema used by the CST table maintained by the data nodes. When the database is created, the table is set up using the schema outlined in the **CurrentSessions.sql** file.

Session control features include:

- User concurrency control—You can enforce concurrency across your entire network by controlling the number of active connections on a per-user, per-cluster basis using the username. You can set a limit on the maximum number of concurrent connections that a user can have and you may use any configured attribute to track instead of just the username.

Subsequently, when the user requests a new connection, the Steel-Belted Radius Carrier servers in the cluster compare the current number of connections to the maximum limit. If a new connection exceeds the limit, the additional connection can be disallowed, or allowed and logged.

- Optional attribute-based concurrency control—If you have a license for the SSR Optional Concurrency and Wholesale Module, you can enforce concurrency across the network by tracking any attribute of a user account, not just the username.
- Cluster session—You can set a limit on the maximum number of concurrent sessions based on your SSR usage requirement.
- Change of Authorization (CoA)/Disconnect Message (DM) processing—You can change the state of authenticated sessions dynamically. In accordance with the Dynamic Authentication RFC 3576, you can:
  - Deploy prepaid scenarios in WLAN integration for GSM/UMTS networks.
  - Place users in legal intercept.
  - Enable online service and billing profile changes.
  - Disconnect or quarantine users for abusing the network.

## RELATED DOCUMENTATION

For information about managing and controlling sessions, see the *SBR Carrier Administration and Configuration Guide*

### Application Support

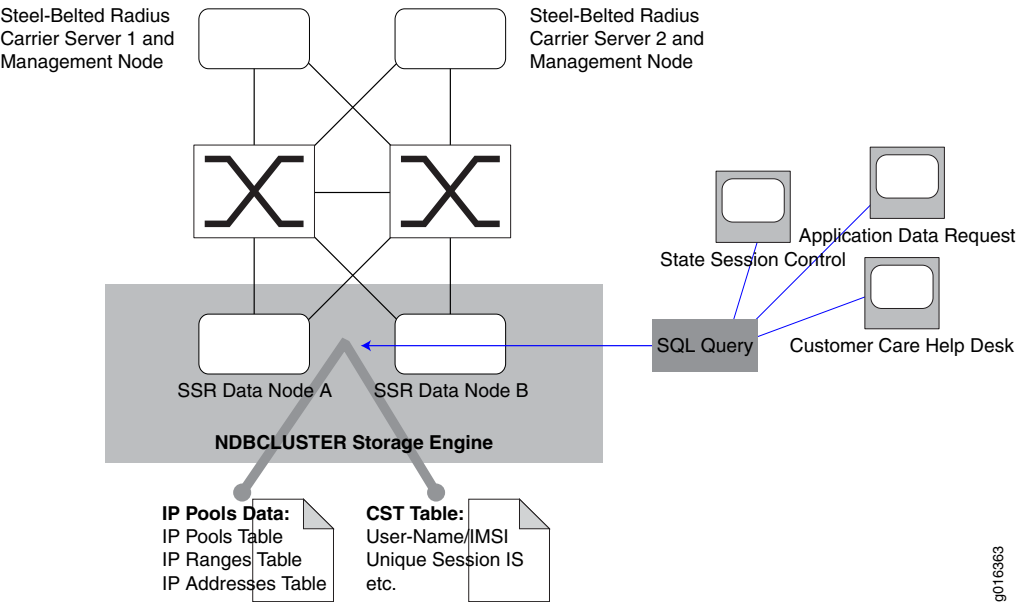
Because the CST is centrally located and can be easily modified, it is often used as the entry point for third party applications that need current session data. This enables Steel-Belted Radius Carrier Session State Register to be the logical source for other applications to query for correlative and verification purposes.

A calling application can make its request through an API as SBR Carrier does, or with a SQL query as shown in [Figure 8 on page 24](#). The SQL query usually takes more processing resources and cycles to complete than an API call, so we recommend against allowing multiple or repeated queries because they can slow down the cluster's primary tasks.

Rows in each table are evenly distributed among a node group's data nodes. A client application can locate a particular record by accessing the unique primary key of any table, which only involves a single node in the cluster. Scans of tables for records meeting broader criteria (such as a search for sessions more than 2 hours old) involves all the data nodes in the primary node group, and takes more processing cycles to complete. We recommend that queries from outside applications be as precise as possible to avoid slowing all operations down.

[Figure 8 on page 24](#) shows the general deployment scenario with usage examples. Communications initiated by applications or systems in the operational support system (OSS) can be formatted as standard SQL queries, shell scripts, or stored procedures, depending on usage requirements.

Figure 8: Database Access by Other Applications



For example, a video gateway might use its own internal identifier (such as an MSISDN, IMSI, or username) for a subscriber, but might not have the subscriber's current IP address. The video gateway can query the Current Sessions Table in real time to retrieve information about an active session.

## Supporting SIM, SIGTRAN, and Other Protocols with Signalware

Steel-Belted Radius Carrier nodes that support SIM depend on Signalware 9 with Service Pack 6.A software being properly installed on the server.

Also, if the EAP-AKA authentication protocol is used by the WiMAX mobility module, Signalware 9 with Service Pack 6.A is required.

**NOTE:** We strongly recommend that Signalware software be installed and configured before Steel-Belted Radius Carrier is installed because the SBR Carrier installation incorporates Signalware settings, if they are available.

Juniper Networks provides an example Signalware installation with:

- The correct settings for simple integration with Steel-Belted Radius Carrier.
- The recommended procedures to configure Signalware 9 to support Steel-Belted Radius Carrier.
- The recommended procedure for migrating configuration settings forward from an existing Signalware installation.

For details, see [“SIGTRAN Support for Steel-Belted Radius Carrier” on page 309](#).

**NOTE:** Because the Signalware communication stack is not supported on Solaris, the SIM authentication module cannot be used on the Solaris platform to communicate with an HLR to process RADIUS requests. However, the module can be used with an HSS by using the RADIUS to Diameter conversion feature.

## Cluster Installation Workflow

The installation workflow follows a list of tasks that must be performed in a specific order to make the Session State Register cluster fully functional. This guide presents groups of related tasks in sequential order.

The high-level view of the installation has several groups of tasks:

1. Planning—Because a SSR cluster involves four or more servers, planning the cluster topology and each server's role and setting is essential. Planning is discussed, and a worksheet you can use to plan your installation are in [“Planning Your Session State Register Cluster” on page 32](#).
2. Before You Begin—The tasks to complete and the information to gather on each server in a cluster before installing software are discussed in detail in [“Before You Install Software” on page 39](#), and in the succeeding chapters that document the optional installation of Signalware that is required to support some features.
3. Software installation—Unpacking the software package, selecting configuration tasks, and providing settings that are used to set up the servers are discussed in detail in [“Installing Session State Register Nodes” on page 94](#).
4. Basic configuration—Setting up configuration files for the SSR database is discussed in [“Customizing the SSR Database Current Sessions Table” on page 144](#).

## Session State Register Cluster Installation Checklist

### Before You Begin

Before you begin to install the Session State Register software and after you read all of this chapter, perform these preliminary tasks or verify that they have been accomplished:

1. Review the Steel-Belted Radius Carrier *Release Notes* for up-to-date information that became available after the Steel-Belted Radius Carrier guides were published.

Download the most recent Release Notes file from:

[https://www.juniper.net/documentation/en\\_US/sbr-carrier8.6.0/information-products/topic-collections/release-notes/sbr-8.6.0.pdf](https://www.juniper.net/documentation/en_US/sbr-carrier8.6.0/information-products/topic-collections/release-notes/sbr-8.6.0.pdf)

If the information in the *Release Notes* differs from the information in any Steel-Belted Radius Carrier guide, follow the *Release Notes*.

2. Read the next chapter, [“Planning Your Session State Register Cluster” on page 32](#), and fill out the provided worksheet so you have all required information ready to use during the installation.

### Migration

3. If you are migrating settings from a SBR/HA 5.5 server installation, put a copy of that installation on the system that you want to host Release 8.6.0.

The recommended method is:

- a. As root on the old server, shut down the SBR daemon.  
Execute: `/radiusdir /sbrd stop`.
- b. Create a gzip archive of the *radiusdir* directory, preserving the file structure.  
Change directories to `/opt/JNPRsbr/` and execute (for example):

```
tar cfE - radius |gzip > 55_radius.tgz
```

- c. As root on the Release 8.6.0 host, create a working directory for the archive:

Execute:

```
mkdir -p /opt/JNPRsbr/PreviousRelease_backup
```

and:

```
cd /opt/JNPRsbr/PreviousRelease_backup
```

- d. Use FTP binary mode to copy the archive file to the Release 8.6.0 host.
- e. Gunzip the archive in the working directory. This does not create a working installation, but serves as a source of files that can be reused by the Release 8.6.0 server.

Execute (using the example file name):

```
gunzip 55_radius.tgz |tar xf -
```

- f. Check the file permissions on the unzipped archive to ensure that the files are writable.

See [“Creating a Copy of Existing SBR Server Release Files for Migration” on page 54](#) for more information.

4. Ensure that all server chassis are physically secure.

To review recommendations on the server’s characteristics, see [“Selecting an Appropriate Server” on page 41](#).

5. Verify root access on all servers.

6. Confirm that all server chassis provide at least the minimum hardware and software requirements.

See [“Meeting System Requirements” on page 41](#).

7. Verify that all servers meet basic network requirements:

- See [“Configuring Multipathing” on page 38](#).
- See [“Verifying Network Connectivity” on page 49](#).
- See [“Verifying Hostname Resolution” on page 50](#).

## CCM

8. Verify the server role in the centralized configuration management (CCM) environment. The installation script can configure a Steel-Belted Radius Carrier server as:

- An autonomous server that does not use CCM.
- The primary server supplying a group of replication servers.
- A replication server.

- In an environment that uses replication servers, the primary server must be configured first to provide a base for replication.

### Migration

- If you are migrating from a SBR/HA 5.5 environment and CCM was in use, you may only migrate the same server type: autonomous to autonomous, primary to primary, and replica to replica.

### LDAP

9. (Optional) If you want the server to interact with a Lightweight Directory Access Protocol (LDAP) directory service agent, verify or record the path to the LDAP library files. (The default path is `/usr/lib`.)

### Oracle

10. (Optional) If you intend the Steel-Belted Radius Carrier server to authenticate using an external database, you need to:

- Verify database compatibility and interoperability.

See ["Setting Up External Database Connectivity \(Optional\)" on page 53](#).

- If the external database is Oracle, the Steel-Belted Radius Carrier server must be configured as a client of the Oracle server.

As you set up the client, collect this information, or if the client is already installed, record this information to use while installing SBR software:

- Record the path to the local Oracle Home directory (for example:  
`/opt/10g/app/oracle/product/10.2.0.3`).
- Record the path to the local Oracle shared library (for example:  
`/opt/10g/app/oracle/product/10.2.0.3/lib`)
- Record the path to the local TNS\_ADMIN (for example:  
`/opt/10g/app/oracle/product/10.2.0.3/network/admin`)

### SIM / WiMAX / SIGTRAN

12. (Required only for these modules) Install and configure Signalware 9 with Service Pack 6.A software, following the product documentation and the guidelines in [“SIGTRAN Support for Steel-Belted Radius Carrier” on page 309](#).

This step is only required for the optional WiMAX mobility module if using EAP-AKA, which requires the optional SIM authentication module. This step is mandatory for the SIM authentication module and any other SIGTRAN applications.

## Installing Session State Register Software

If you have completed all items on the “Before You Begin” list (for each item for each server in the cluster), you are ready to install the software. The servers must be installed in a particular order, so you perform this portion of the list four times for a Starter Kit.

The order of installation for a four-server Starter Kit cluster is:

1. First SBR Carrier/management node host. This server is also the primary server if CCM is used because it must be installed before the replica nodes.
2. Second SBR Carrier/management node host.
3. First data node host.
4. Second data node host.

Repeat these key steps on each server:

1. Unpack the Session State Register software.

See [“Unpacking Session State Register Software” on page 95](#).

If you are not familiar with UNIX package management commands, see [“Package Management Commands” on page 95](#).

2. Run the Session State Register **configure** script to set up or read in cluster and server configuration files.
3. Restart the Session State Register **configure** script.

When the script runs the second time, you are prompted to enter more information that you recorded in [“Planning Your Session State Register Cluster” on page 32](#) and [“Before You Install Software” on page 39](#).

## Basic Configuration

After the Session State Register software is on all servers in the cluster, you can ensure that the base software runs properly and then perform basic configuration steps:

1. Start the cluster.

See [“Initial Cluster Startup Procedure” on page 132.](#)

2. Launch the Web GUI.

Use an HTML browser to address one of the Steel-Belted Radius Carrier servers. Download the Administrator client application and log in as root. Also confirm that the links to the documentation work. See [“Launching Web GUI” on page 138.](#)

Verify that you can download and run the application both from the local host and from a workstation on the network.

3. Configure the server.

For a new installation, configuring the environment includes basic low-level configuration tasks such as:

- Configuring communications between any additional devices and the Steel-Belted Radius Carrier servers.
- Adding additional users to the server.

See [“Configuring the Server” on page 140.](#)

4. Configure SNMP. If you elected not to configure the Steel-Belted Radius Carrier SNMP agent during software installation and other agents also run on the server (such as the Solaris operating system agent), you must adjust the ports the agents use so multiple agents do not contend for the same resource.

See [“Configuring SNMP” on page 141.](#)

5. Set up shared SSR IP address pools and a local emergency pool. See [“Setting Up IP Address Pools” on page 142.](#)



# 2

PART

## Preparing for a Steel-Belted Radius Carrier Installation

---

Planning Your Session State Register Cluster | 32

Before You Install Software | 39

Migrating from Previous SBR Releases | 56

---

# Planning Your Session State Register Cluster

IN THIS CHAPTER

- Planning the Cluster Topology | 32
- Cluster Naming Worksheets | 34
- Renaming Node Hosts | 38
- Configuring Multipathing | 38


This chapter documents the planning and preparatory tasks that are required to set up a Session State Register cluster. *If you are installing a standalone SBR Carrier server, skip this chapter and proceed to Chapter 4, “Before You Install Software” on page 39.*

These topics are in the chapter:

## Planning the Cluster Topology

The topology of all clusters begins with the four machines required to implement the SSR Starter Kit: two machines that each host both an SBR node and an SSR management node and two machines that each host one SSR data node.

Figure 9: Basic Steel-Belted Radius Carrier Session State Register Starter Kit Cluster



## Naming the Cluster and Its Machines

To determine the machine names:

1. Name the cluster.

The cluster name must:

- Begin with a letter (an alphabetic character).
- Be made up of alphanumeric characters, underbars, and dashes.
- Be a single string without any spaces.

2. Name the nodes. You must set the node name to be the same as the hostname associated with the IP address of the primary network interface. You can test the name with the **uname -n** command.

It is not required, but you can simplify machine identification if you adopt a convention that builds on the node identification numbers (node ID) that are used by the SSR processes and incorporate them into the server name.

Guidelines for the assignment of node IDs:

- 0 is for internal use.
- 1–48 are for clustered data nodes.
- 51–59 are for clustered management nodes.
- 61–69 are for clustered management nodes (a function of base node ID + 10).
- 70–99 are reserved for future use.
- 100–149 are for clustered SBRC nodes.
- 150–255 are reserved for future use.
- 256 and higher are not supported.

**NOTE:** If striping is enabled, you are restricted to the range 1–N for data node IDs, where N is the total number of data nodes in the cluster.

### ***Starter Kit Cluster Naming Example***

Table 8 on page 34 shows a starter kit cluster naming example. The cluster and node names in Table 8 on page 34 are used in subsequent chapters to create a working cluster example. Because multipathing is required, only one IP address is used for each machine; in this example, because the multipath addresses are assigned, each server's IP address is set to echo each server's primary node ID number.


**NOTE:** The node IDs used in Table 8 on page 34 and Figure 10 on page 34, are examples only and do not accurately reflect the assignment of node IDs.

Table 8: Example Cluster Naming Worksheet

Product Name	Cluster Name	Hosted Nodes	Node Name	Multipath IP Address	Node ID Numbers
SSR Starter Kit	MyCluster	SBR and Mgt.	MyNode_1	192.168.0.1	ndb_mgmd =1
					mysqld =6
					radiusd =41
		SBR and Mgt.	MyNode_2	192.168.0.2	ndb_mgmd =2
					mysqld =7
					radiusd =42
		Data	MyNode_3	192.168.0.18	ndbd =18
		Data	MyNode_4	192.168.0.19	ndbd =19

Implementing the worksheet in [Table 8 on page 34](#) results in the cluster setup shown in [Figure 10 on page 34](#).

Figure 10: The Starter Kit Cluster



## Cluster Naming Worksheets

Use [Table 9 on page 35](#) to plan your SSR Starter Kit cluster. The table includes entries for the four Starter Kit systems. You can ignore the Node ID Numbers column for now; it is explained and used later. If you use more nodes, see [Table 10 on page 35](#) for an Expansion Kit worksheet.

**NOTE:** The IP address for the SBR node process and management node process in a SM node must be the same for both node processes, otherwise the **configure** script forces the SBR node process address to the IP address of the management node process.

**Table 9: Starter Kit Cluster, Machine, and Interface Naming Worksheet**

Product Name	Cluster Name	Hosted Nodes	Node Name	Multipath IP Address	Node ID Numbers
SSR Starter Kit		SBR and Mgt.		. . .	ndb_mgmd =
					mysqld =
					radiusd =
		SBR and Mgt.		. . .	ndb_mgmd =
					mysqld =
					radiusd =
		Data		. . .	ndbd =
		Data		. . .	ndbd =

### Expansion Kit and Additional Node Worksheet

You can add two data nodes (a total of four), one management node (a total of three), and up to 18 additional SBR Carrier nodes (a total of 20) to a cluster. Use the worksheet in [Table 10 on page 35](#) to plan out the addition of expansion kits and additional nodes.

**Table 10: Additional Node Machine, and Interface Naming Worksheet**

Product Name	Cluster Name	Nodes	Node Name	Multipath IP Address	Node ID Numbers
First Expansion Kit		Data			ndbd
		Data			ndbd

Table 10: Additional Node Machine, and Interface Naming Worksheet *(continued)*

Product Name	Cluster Name	Nodes	Node Name	Multipath IP Address	Node ID Numbers
Third Management Node		Management			ndb_mgmd



## Renaming Node Hosts

If you change the names of any hosts, set the corresponding node names to be the same as the hostnames. You can use the **uname -n** command to test the hostname.

## Configuring Multipathing

To ensure a redundant installation with redundant network interfaces, and to work with the installation scripts, all servers in the cluster need to have multipathing configured. We recommend static IP addresses.

The general steps required to set up IP multipathing on a Solaris machine are:

1. Ensure that the system detects both network interface cards and they have different MAC addresses.
2. Group the network interface cards.
3. Add a test IP address to the first network interface card.
4. Add a test IP address to the second network interface card.
5. Edit each server's **/etc/hosts** file.
6. Change the **hostname.\*** files to keep these settings after a reboot.

**BEST PRACTICE:** When activating IP multipathing on a server, a total of three IP addresses are used. If the private address space is a Class C space, only 254 usable addresses are usually available.

To keep the IP address space uncluttered, a convention for IP assignments has evolved—assign the main failover IP to be in the range 1–54, the first test address 100 higher than the main address, and the second test address 200 higher than the main address. So for a main address of 192.168.0.41, the first test address is 192.168.0.141 and the second 192.168.0.241.

Create DNS entries for the test addresses although the addresses are not used. The convention for this is to use the hostname of the server with **-test1** or **-test2** appended to the end.

Going back to the sample worksheet ([Table 8 on page 34](#)), hostname **MyNode\_1** uses a multipath IP of 192.168.0.1 and its interfaces are **MyNode\_1-test1** at 192.168.0.101 and **MyNode\_1-test2** at 192.168.0.201.



# Before You Install Software

## IN THIS CHAPTER

- Reviewing the Release Notes | 39
- Determining the Server's Centralized Configuration Management Role | 40
- Selecting an Appropriate Server | 41
- Meeting System Requirements | 41
- Verifying Root Access | 49
- Verifying Network Requirements | 49
- Creating Required Users and Groups | 50
- Creating Share Directories | 52
- Setting Up External Database Connectivity (Optional) | 53
- Creating a Copy of Existing SBR Server Release Files for Migration | 54
- Installing the SIGTRAN Interface (Optional) | 55

This chapter presents requirements that you must check and tasks that you must perform before installing any Steel-Belted Radius (SBR) Carrier software. This chapter and its topics apply to all SBR Carrier servers, both standalone servers and SBR Carrier node hosts in a Session State Register cluster, and to servers in a cluster that do not run SBR Carrier software themselves—management and data node hosts.

These topics are in the chapter:

To understand more about SBR Carrier Licenses, see, *Obtaining License Keys*. Please refer to the *Licensing Guide* for general information about License Management. Please refer to the product [Data Sheets](#) for further details, or contact your Juniper Account Team or Juniper Partner.

## Reviewing the Release Notes

The Steel-Belted Radius Carrier *Release Notes* contain important late-breaking information, such as known software problems. Please review the release notes delivered with your Steel-Belted Radius Carrier software before you install SBR Carrier to ensure you are informed about important information not found elsewhere.

Download the Release Notes from:

[https://www.juniper.net/documentation/en\\_US/sbr-carrier8.6.0/information-products/topic-collections/release-notes/sbrc-8.6.0.pdf](https://www.juniper.net/documentation/en_US/sbr-carrier8.6.0/information-products/topic-collections/release-notes/sbrc-8.6.0.pdf)

## Determining the Server's Centralized Configuration Management Role

Steel-Belted Radius Carrier software includes an optional feature called centralized configuration management (CCM). CCM coordinates Steel-Belted Radius Carrier server settings between a primary RADIUS server and one or more replica RADIUS servers. It copies critical configuration files from the primary to the replicas, so it keeps multiple SBR Carrier servers operating the same way.

Although we recommend CCM for multiserver installations, its use is optional. If necessary, an environment can contain any number of unique server configurations. A server that does not use CCM is an *autonomous* server. In a CCM group one server is the *primary* server and the others are *replicas*.

If you are installing in an environment that uses CCM, make sure that the primary server is fully functional before installing any replicas—the replica servers need to communicate with the primary server during installation and when they start to acquire updated configuration files.

**TIP:** In this context, *primary* and *replica* refer to CCM. *Primary* and *replica* are used in another context when the topic is data management within a Session State Register cluster. CCM is a separate tool and process independent of SSR, but it is often used in SSR environments to keep the SBR Carrier nodes operating identically.

For more information about CCM, see *Configuring Replication* in the *SBR Carrier Administration and Configuration Guide*.

**TIP:** In the context of data management within a Session State Register cluster, these terms refer to the data that data nodes manage. For more information, see [“SSR Cluster Concepts and Terminology” on page 10](#)

**NOTE:** If you are migrating from an earlier SBR release to Release 8.6.0, both servers must be the same CCM role. For example, a Release 6.1 primary server can be migrated forward to a Release 8.6.0 primary server. Or an autonomous server can be used to set up a new autonomous server. But cross-type migration is not supported.

## Selecting an Appropriate Server

Select a host with these properties to run the Steel-Belted Radius Carrier server software:

- A secure physical location such as a locked server room.
- Root access on the host limited to the system administrator. Restrict login access to the Steel-Belted Radius Carrier server to system administrators and others who need it. Limit the server to no (or few) user accounts.
- Adequate memory and disk space—See [“Meeting System Requirements” on page 41](#) for information about hardware and software requirements.
- Administrative interface not accessible from outside your network. If your Steel-Belted Radius Carrier server has one network connection, limit access to the ports that it uses for configuration and administration.

If your Steel-Belted Radius Carrier server has more than one network connection, use an administrative network that is physically separate from other networks to configure and administer the server.

- Does not provide public network services such as FTP or HTTP. Although FTP is often used during installation, disable it afterwards. Turning FTP off both frees some resources and closes a potential entry point to malicious attacks. (Steel-Belted Radius Carrier implements its own HTTP service on a different socket port than generic web servers use.)
- A shared secret protects all communications to and from the server, including session keys for wireless data encryption. Configure shared secrets that are long enough and random enough to resist attack. Avoid using the same shared secret throughout your network.

Specific hardware requirements are discussed in [“Meeting System Requirements” on page 41](#).

## Meeting System Requirements

This section describes the hardware and software requirements for running Steel-Belted Radius Carrier on the following server and operating system combinations:

- Oracle Sparc server and Solaris operating system
- x86 server and Linux operating system

### Standalone SBR Carrier Server Hardware

These basic specifications may be used for any standalone Steel-Belted Radius Carrier server—one that does not participate in a Session State Register cluster.

Session State Register servers have additional system requirements, such as dual Gigabit Ethernet NICs (to provide redundant communication links). These additional requirements are discussed in [“Supported SBR Carrier SSR Cluster Configurations”](#) on page 16.

Table 11 on page 42 lists the hardware requirements for 64-bit standalone SBR Carrier server.

**NOTE:** These hardware requirements are also applicable to VMs hosting standalone SBR Carrier in virtualized environments. Overprovisioning is not recommended.

Table 11: 64-Bit Version of Standalone SBR Carrier Server Hardware Configurations

Server	RAM	Solaris CPU	Linux CPU	Free Disk Space
64-Bit Standalone SBR Carrier Server (Minimum Configuration)	8 GB RAM. <sup>1</sup>	Two-CPU M5000 system or better, running at 2.6 GHz.	Xeon 2-core at 2.0 GHz supported for testing.  Xeon 4-core or 2x2-core at 2.4 GHz supported for limited performance production use.	At least 8 GB of local hard disk space (not NFS).

Table 11: 64-Bit Version of Standalone SBR Carrier Server Hardware Configurations (*continued*)

Server	RAM	Solaris CPU	Linux CPU	Free Disk Space
<b>64-Bit Standalone SBR Carrier Server (Recommended Configuration)</b>	<p>16 GB RAM or more.</p> <p>If Mitel's Signalware communication stack is used, for systems processing a heavier-than-normal load (for instance, with WiMAX or additional session licenses), or when configuring a large number of threads (for example, to accommodate high latency of proxy targets), more memory produces better performance.</p>	SPARC T5 or better, running at 3.6 GHz.	<p>Xeon 2x6-core at 3.6 GHz or higher.</p> <p>Higher core counts with lower processing speeds may not improve throughput in certain cases. See the <i>Performance, Planning, and Tuning Guide</i> for more information.</p>	At least 16 GB of local hard disk space (not NFS).

1. The memory size can be increased based on the subscriber count.

## Session State Register Host Hardware

### *SBR Carrier and Management Node Hosts*

[Table 12 on page 44](#) lists the hardware requirements for Session State Register cluster SBR Carrier and management node hosts for 64-bit SBR Carrier server.

**Table 12: Session State Register SBRC and Management Node Host Hardware Configurations for 64-Bit SBR Carrier**

Server	RAM	Solaris CPU	Linux CPU	Free Disk Space	Network Interfaces
<b>64-Bit SBRC and/or Management Node Host (Minimum Configuration)</b>	8 GB RAM.	SPARC T5 or better, running at 3.6 GHz.	Xeon 2-core at 2 GHz supported for testing and M-only node.  Xeon 4-core or 2x2-core supported for production for a single SM node.	At least 8 GB of local hard disk space (not NFS).	Two physical interfaces on a 100 Base-T network.  Multipath configuration is required.
<b>64-Bit SBRC and/or Management Node Host (Recommended Configuration)</b>	16 GB RAM or more.  If Mitel's Signalware communication stack is used, for systems processing a heavier-than-normal load (for instance, with WiMAX or additional session licenses), or when configuring a large number of threads (for example, to accommodate high latency of proxy targets), more memory produces better performance.	SPARC S7 or better, running at 4.2 GHz.	Xeon 2x4-core at 3.6 GHz or higher.  Higher core counts do not improve performance. See the <i>Performance, Planning, and Tuning Guide</i> for more information.	At least 16 GB of local hard disk space (not NFS).	Two physical interfaces on a Gigabit Ethernet network.  Multipath configuration is required.

### Data Node Hosts

Table 13 on page 45 lists the hardware requirements for Session State Register data node hosts.

All data node hosts in a cluster must have the same configuration. Because they collaborate to keep a shared database in virtual shared memory, the processing power, RAM, and communications capability of all the machines need to be very similar.

**NOTE:** This free disk space shown in Table 13 on page 45 must be available specifically to the /opt file system for installation of the SSR software.

Table 13: Session State Register Data Node Host Hardware Configurations

Server	RAM	Solaris CPU	Linux CPU	Free Disk Space	Network Interfaces
<b>Data Node Host (Minimum Configuration)</b>	10 GB RAM.	SPARC T5 or better, running at 3.6 GHz.	Xeon 2-core at 2.0 GHz supported for testing.  Xeon 4-core or 2x2-core at 2.4 GHz supported for production.	The local disk space requirement is related to the amount of RAM assigned to the data node. To calculate the minimum requirement for the amount of RAM on the node, use the formula: $(\text{RAM} - 4 \text{ GB}) * 12$ .	Two physical interfaces on a 100 Base-T network.  Multipath configuration is required.
<b>Data Node Host (Recommended Configuration)</b>	More than 10 GB RAM.  More than the minimum of 10 GB of RAM supports more connections because more of the SSR database can be held in memory. More database in memory may translate into faster processing because disk operations are minimized.  <b>NOTE:</b> In particular, you have to manually increase the memory set to the <b>DataMemory</b> parameter in <b>config.ini</b> if you want the SSR to store more CST data.	SPARC S7 or better, running at 4.2 GHz.	Xeon 2X4-core at 3.6 GHz or higher for highest throughput.	For example, a data node with 16 GB of RAM requires a minimum of $(16 \text{ GB} - 4 \text{ GB}) * 12$ , or 144 GB of local disk storage space.	Two physical interfaces on a Gigabit Ethernet network.  Multipath configuration is required.

**Checking Free Disk Space**

Use the **df** command to verify the amount of free disk space on a potential server.

Execute:

```
df -hk
```

Figure 11: Checking Free Disk Space with the df Command

```

# df -hk
Filesystem            size  used  avail capacity  Mounted on
/dev/dsk/c0t0d0s0      1.9G   145M   1.7G      8%      /
/devices              OK      OK      OK      0%    /devices
ctfs                  OK      OK      OK      0%    /system/contract
proc                  OK      OK      OK      0%    /proc
mnttab                OK      OK      OK      0%    /etc/mnttab
swap                  9.1G   1.4M   9.1G      1%    /etc/svc/volatile
objfs                 OK      OK      OK      0%    /system/object
/dev/dsk/c0t0d0s4      7.7G   3.9G   3.7G     52%    /usr
/platform/sun4u-us3/lib/libc_psr/libc_psr_hwcapl.so.1
/platform/sun4u-us3/lib/libc_psr/libc_psr.so.1
/platform/sun4u-us3/lib/sparcv9/libc_psr/libc_psr_hwcapl.so.1
/platform/sun4u-us3/lib/sparcv9/libc_psr/libc_psr.so.1
fd                    OK      OK      OK      0%    /dev/fd
/dev/dsk/c0t0d0s3      9.6G   1.3G   8.2G     14%    /var
swap                  9.1G   216K   9.1G      1%    /tmp
swap                  9.1G    64K   9.1G      1%    /var/run
/dev/dsk/c0t0d0s5      42G    5.1G   37G     13%    /opt
/dev/dsk/c0t0d0s6      1.9G   2.0M   1.8G      1%    /reserved
/dev/dsk/c0t1d0s0      63G   1.2G   62G      2%    /export/local
wf-filer3:/vol/voll/homes/akashyap
wf-filer3-wf:/vol/homes/homes/ecoyne
sdc-filer2-8:/vol/bread_home3/home8/jasyed

```

## Software

The Steel-Belted Radius Carrier software can be installed on both Solaris and Linux operating systems.

**NOTE:** You cannot run multiple instances of SBR Carrier on Solaris and Linux platforms.

Make sure that nss-util, nss, nspr, gcc, openldap, Kerberos, cyrus-sasl, and zlib libraries are installed on your system before installing the SBR Carrier software. The libraries are normally available and installed with the OS base bundle. SBR Carrier supports the package versions of preceding libraries that are provided with Solaris 11.3.36.10.0 or later and RHEL 7.3 or later.

We recommend you to update the OS regularly for security reasons. Any questions concerning vulnerabilities of libraries that are not distributed by Juniper Networks should be addressed to the OS vendor (Red Hat or Oracle).



## Solaris

SBR Carrier has been qualified on and supports Oracle Solaris 11.3.36.10, 11.3.36.20, and 11.4.25.0.1.75.3.

**NOTE:** Because the Signalware communication stack is not supported on Solaris, the SIM authentication module cannot be used on the Solaris platform to communicate with an HLR to process RADIUS requests. However, the module can be used with an HSS by using the RADIUS to Diameter conversion feature.

### Special Notes on Solaris Platform Support

You need to consider the following special notes while using SBR Carrier on a Solaris platform:

- To use the command-line tools (snmpget, scriptcheck, ldapsearch, and so on) for 64-bit Solaris, first set the library path environment variable and then run the command.

To set the library path environment variable for 64-bit Solaris, execute the following strings:

```
LD_LIBRARY_PATH_64=radiusdir:radiusdir/system/lib/:radiusdir
/openldap:/usr/sfw/lib/sparcv9:/usr/lib/mps/sparcv9:/usr/sbin/lib/sparcv9:/lib/sparcv9:/
usr/lib/sparcv9:/usr/openwin/lib/sparcv9:/usr/dt/lib/sparcv9:/usr/proc/lib:/
usr/local/lib/sparcv9:/opt/sfw/lib/sparcv9:/usr/ccs/lib/sparcv9:/usr/ucblib/sparcv9:
```

## Linux

SBR Carrier supports Red Hat Enterprise Linux 7.3, 7.4, 7.5, 7.6, 7.7, 7.8, and 8.1 on Intel (Xeon) hardware.

**NOTE:** SBR Carrier does not support RHEL 6.x, 7.0, 7.1, and 7.2.

By default, SCTP module is blacklisted in RHEL8. Refer the *Blacklisting SCTP module by default in RHEL 8* section of the [RHEL8 release notes](#) to enable the SCTP module for Diameter Feature.

The SBR standalone and cluster configurations are supported on the 64-bit version of Linux operating system.

**NOTE:** SBR automatically updates the required shared packages from the Red Hat repository. A current Red Hat Enterprise Linux support subscription and an access to the Red Hat repository or a local repository are required to get the latest updates to shared packages. See the appropriate Red Hat administration manuals for more information.

### Special Notes on Linux Platform Support

You need to consider the following special notes while using SBR Carrier on a Linux platform.

- The authGateway and GWrelay processes must be restarted whenever SBR restarts. This is applicable only on Linux.
- SBR installation depends on the binary calculator (bc) command-line tool because the **AddPools.sh** script fails to run if the “bc” tool is not installed.

Table 14 on page 48 lists the features that are limited to support on specific operating systems.

**Table 14: Features and Their Supported OS**

Feature	64-bit Linux Support?	64-bit Solaris Support?	VM Support?
authGateway application	Yes	Yes	Yes
SIM authentication module and GWrelay application	Yes	Yes	Yes
SSR cluster	Yes	Yes	No <sup>*</sup>
Customer-written SDK plug-ins	Yes	Yes	Yes

<sup>\*</sup> SSR cluster in virtualized environments is not officially supported. Juniper Networks may still provide support for known issues and for those where you can demonstrate the issue exists on the native OS.

Subject to the preceding limitation, SBR Carrier supports virtualization on Linux, VMware hypervisor, Kernel-based Virtual Machine (KVM) hypervisor, and logical domains on Solaris. SBR Carrier has been tested with VMware ESXi 5.1, 5.5, 6.0, and 6.5 versions and KVM hypervisor on a RHEL 7.3 machine. For more information on planning and tuning the performance of SBR Carrier running on the Linux and Solaris operating systems, see the *Performance, Planning, and Tuning Guide*.

### Perl

Oracle Solaris 11 is shipped with Perl 5.12.5, and Steel-Belted Radius Carrier has been tested with this version. Multiple Perl installations in discrete directories are supported, but attempting to use other versions of Perl with SBR Carrier may cause problems.

After installation, be sure that the first line of the **radiusd** script specifies the correct path to the Perl 5.12.5 executable. For example, if Perl 5.12.5 is installed in the default location, `/usr/bin/perl`, then the first line of the **radiusd** script must be:

```
#!/usr/bin/perl
```

### LDAP Plug-in

The LDAP plug-in requires SASL, which is included with the SBR Carrier package only for Solaris versions and not for Linux versions. For a Linux machine, you must ensure that you have the SASL package installed before starting SBR.

## Verifying Root Access

You must have access to the **root** account on all machines to install Steel-Belted Radius Carrier software.

## Verifying Network Requirements

Basic network connectivity is required to complete the tasks in this chapter. Although it is not a strict requirement, we recommend that a production site use separate networks for RADIUS traffic, SBR administration and SNMP, and SSR traffic.

### Verifying Network Connectivity

Use the ping command both to and from another computer to verify that the server on which you are going to install Steel-Belted Radius Carrier can communicate with other devices, such as remote access servers, database servers, DHCP servers, DNS servers, and management workstations on your network, using TCP/IP.

- If you are going to use CCM, make sure that the primary server and all replica servers can communicate with each other.
- If you are going to use Session State Register, make sure that all node hosts can communicate with all other node hosts.

For example:

**ping 192.168.12.54**

```
Reply from 192.168.12.54: bytes=32 time=7ms TTL=255
Reply from 192.168.12.54: bytes=32 time=7ms TTL=255
Reply from 192.168.12.54: bytes=32 time=7ms TTL=255
Reply from 192.168.12.54: bytes=32 time=7ms TTL=255
```

If the ping command fails, verify that the IP address of the Steel-Belted Radius Carrier host is correct, that is operational, and that all routers between your server and the remote host are operational.

## Verifying Hostname Resolution

The server on which you are going to install Steel-Belted Radius Carrier must have a stable, accessible IP address that is mapped in `/etc/hosts` or the Domain Name System (DNS) server to a resolvable hostname.

**NOTE:** If you are using installing Session State Register hosts, which must have two network interface cards to provide redundant communications, be sure to test both unless multipathing has already been set up.

To verify that the server has a resolvable hostname, use the ping command with the server's hostname. For example:

**ping foo.juniper.net**

```
Pinging foo.juniper.net [192.168.12.21] with 32 bytes of data:
Reply from 192.168.12.21: bytes=32 time=7ms TTL=255
Reply from 192.168.12.21: bytes=32 time=7ms TTL=255
Reply from 192.168.12.21: bytes=32 time=7ms TTL=255
Reply from 192.168.12.21: bytes=32 time=7ms TTL=255
```

## Creating Required Users and Groups

All Steel-Belted Radius Carrier servers and all nodes in a Session State Register cluster require a **hadm** user account that is part of the **hadmg** group.

**NOTE:** Make sure the **hadm user id** and **hadmg group id** are the same on all machines in the cluster.

- Ensure the home directory for the **hadm** user is `/opt/JNPRhadm`
- Ensure the shell is `/bin/bash`.

To determine whether a **hadm** user account and **hadmg** group already exist on the server:

1. As root, execute **getent group hadmg** and look for a hadmg listing.

Execute:

**getent group hadmg**

If the group exists, a message similar to this example is displayed:

```
root@crispix:/opt> getent group hadmg
```

```
hadmg: :65536:
```

If the group exists, and was used for a purpose other than SBR Carrier, or if the group ID number is not the one you intend to use, then you need to terminate all references to the group (which is beyond the scope of this document), and delete the group so that it can be re-created and reassigned to SBR Carrier. If the group ID number is correct, then you do not need to delete the group.

Execute:

**groupdel hadmg**

2. As root, execute **getent passwd hadm** and look for a hadm listing.

Execute:

**getent passwd hadm**

If the hadm user exists, a message similar to this example is displayed:

```
root@crispix:/opt> getent passwd hadm
```

```
hadm:x:16761:65536:SBR software:/opt/JNPRhadm:/bin/bash
```

If the **hadm** user exists and was used for a purpose other than SBR Carrier, or if the user ID number is not the one you intend to use, then you need to terminate all references to the user (which is beyond the scope of this document), and delete the user so that it can be re-created and reassigned to SBR Carrier. If the user ID number is correct, then you do not need to delete the user. In any case, be sure to back up the existing home directory first, for example, by renaming it.

Execute:

**mv /opt/hadm /opt/hadm.old**

and

**userdel hadm**

3. Verify that **hadm** and **hadmg** have read and write permissions to the **/tmp** directory.

If the **hadm** user account and **hadmg** group do not already exist on the server:

1. Create the **hadm** account and **hadmg** group.

- a. As root, create the **hadmg** group account with the intended group ID number. According to Solaris convention, specify a group ID number in the range 100–65535 inclusive; for example, 7772.

Execute:

```
groupadd -g 7772 hadmg
```

- b. Create the **hadm** user account with the intended user ID number and **hadmg** as the primary group account. According to Solaris convention, specify a user ID number in the range 100–65535 inclusive; for example, 9969.

Execute (all on one line):

```
useradd -u 9969 -g 7772 -c "SBR Software" -d /opt/JNPRhadm -s /bin/bash hadm
```

- c. Confirm that the group and user were created.

Execute:

```
getent passwd hadm
```

```
getent group hadmg
```

- d. Set the **hadm** user account password.

Execute:

```
passwd hadm
```

The **hadm** home directory is created by the **configure** script in a later task.

2. Verify that **hadm** and **hadmg** have read and write permissions to the **/tmp** directory.

## Creating Share Directories

A share directory containing some subdirectories can be used to distribute master files between like machines. If you choose to create a share directory, it is used among different nodes in the cluster, and to transfer configuration files from previous releases of SBR server.

To create the share directories:

1. As root, create the share directory.

Execute:

```
mkdir -p /opt/JNPRshare
```

```
chmod 775 /opt/JNPRshare
```

2. Assign the ownership to **root** and **hadmg**.

Execute:

```
chown root:hadmg /opt/JNPRshare
```

3. Move to the share directory.

Execute:

```
cd /opt/JNPRshare
```

4. Create an **/install** subdirectory.

Execute:

```
mkdir -p install  
chmod 775 install
```

5. Assign the ownership to **hadm** and **hadmg**.

Execute:

```
chown hadm:hadmg install
```

6. Move to the **/install** directory.

Execute:

```
cd install
```

7. Create a hostname-specific or SSR cluster-specific subdirectory.

Execute:

```
mkdir -p < hostname >|< cluster-name >  
chmod 775 < hostname >|< cluster-name >
```

- If you are installing on a machine to be a standalone SBR Carrier server, use the machine's hostname.
- If you are installing on a machine to be part of a Session State Register cluster, use the cluster name.  
If you are not sure what it is, see ["Naming the Cluster and Its Machines" on page 32](#).

8. Assign the ownership to **hadm** and **hadmg**.

Execute:

```
chown hadm:hadmg < hostname > or < cluster-name >, whichever you used.
```

## Setting Up External Database Connectivity (Optional)

If the Steel-Belted Radius Carrier server or hosts support use an external database for authentication, verify that it is a supported database, and if necessary, set up the server:

- On a Solaris machine, Oracle 10, 11, and 12 are supported; versions 10.2.0, 11.2.0, and 12.1.0.2 are recommended.

On a Linux machine, Oracle 11 and 12 are supported; versions 11.2.0 and 12.2.0 are recommended.

The SBR Carrier server must be configured as a client of the Oracle server before installing SBR Carrier, and the Oracle server location is used during installation.

Follow the directions in the Oracle documentation to install and configure the client, and record this information to use while installing SBR Carrier software:

1. Record the path to the local Oracle Home directory (for example:  
`/opt/10g/app/oracle/product/10.2.0.3`)
  2. Record the path to the local Oracle shared library (for example:  
`/opt/10g/app/oracle/product/10.2.0.3/lib`)
  3. Record the path to the local TNS\_ADMIN (for example:  
`/opt/10g/app/oracle/product/10.2.0.3/network/admin`)
- The JDBC plug-in has been tested with Oracle on Solaris and the JDBC plug-in for MySQL.

## Creating a Copy of Existing SBR Server Release Files for Migration

If you are already operating standalone servers that run SBR Carrier Release 6.0, Release 6.1, or Release 7.0, and you want to carry the existing configuration forward during the Release 8.6.0 installation, you need to transfer a copy of the existing installation to the machine that you want to host SBR Carrier. This is a non-functional copy that the installation scripts can access in order to selectively import some configuration files. The exact files that can be reused vary from release to release. For details on standalone migration, see [“Migrating from Previous SBR Releases” on page 56](#).

To create a copy:

1. As root on the existing server, shut down the SBR daemon.

Execute:

```
/radiusdir /sbrd stop.
```

2. Create a gzip archive of the old *radiusdir* directory, preserving the file structure.

Change directories to `/opt/JNPRsbr/` and execute:

```
tar cfE - radius |gzip > radius.tgz
```

3. As root on the Release 8.6.0 server or host, create a working directory for the archive:

Execute:

```
mkdir -p /opt/JNPRsbr/PreviousInstall
```

and:

```
cd /opt/JNPRsbr/PreviousInstall
```



**NOTE:** If you are using CCM, remember that you can only migrate server configurations between machines of the same type—primary to primary and replica to replica. So if you are migrating both a primary server and one or more replicas, you need to create a backup directory for each because they contain different types of information.

- The backup of the primary is used when you install the new primary server
- The backup of the replica is used when you install the new replicas.

4. Copy the archive file to the appropriate Release 8.6.0 host.
5. Gunzip the archive in the working directory. This does not create a working installation, but serves as a source of files that can be reused by the Release 8.6.0 server.

Execute:

**gunzip radius.tgz |tar xf -**

6. Check the file permissions on the unzipped archive to ensure that the files are writable.

Remember the path to the archive directory. During installation, the Release 8.6.0 installation script prompts you to supply the location so some configuration files can be migrated into the new installation.

## Installing the SIGTRAN Interface (Optional)

If the SIM authentication plug-in is used, Signalware 9 with Service Pack 6.C needs to be installed in the server before you install SBR Carrier software.

See [“SIGTRAN Support for Steel-Belted Radius Carrier” on page 309](#) for information and procedures to install and configure Signalware configuration files to support SBR Carrier.

# Migrating from Previous SBR Releases

## IN THIS CHAPTER

- Migration Overview | 56
- Files from Earlier Releases That Require Manual Editing | 58
- Migrating from SBR Release 6.X | 59
- Migrating from SBR Carrier Release 7.0 | 59
- Migrating from SBR Carrier Release 7.4.0 | 60
- Migration and New Installations of SBR Carrier with WiMAX | 60
- Historic Files | 61

This chapter documents the general—and optional—procedure to migrate configuration files from previous SBR releases into a Steel-Belted Radius Carrier Release environment. This chapter applies only to standalone servers, not to servers in a Session State Register environment and only if you imported configuration files from a previous release during installation (discussed in [“Running the Steel-Belted Radius Carrier Configure Script” on page 68](#)).

This chapter contains these topics:

## Migration Overview

You can use the configuration script to move a number of files from selected previous releases from the old platform into the new Release environment when installing Steel-Belted Radius Carrier. The corresponding Release files are also loaded on the system, but are not activated. You are responsible for merging new settings from Release configuration files into the working (old) configuration files. To support new features, SBR Carrier uses default values for any new settings that have not been merged into the working configuration files.

**NOTE:** During migration the logs are written to the configure.log file.

## Supported Releases

Configuration files from these releases of SBR server installations may be carried forward into Release 8.6.0:

- SBR Release 6.0 and Release 6.1
- SBR Carrier Release 7.x.x
- SBR Carrier Release 8.0.0, 8.1.0, 8.2.0, 8.3.0, 8.4.0, 8.4.1, and 8.5.0

## Files That May Be Migrated

Depending upon the source release, there may be some limitations or recommendations on copying and merging certain other binaries, configuration files, and data. These restrictions are listed in the release-specific sections that follow. In most cases, though, files of these types are rolled forward by the **configure** script:

- \*.acc
- \*.att
- \*.aut
- \*.conf
- \*.ctrl
- \*.dat
- \*.dci (See [“Manual Migration of Dictionaries” on page 59.](#))
- \*.dcm (See [“Manual Migration of Dictionaries” on page 59.](#))
- \*.dct (See [“Manual Migration of Dictionaries” on page 59.](#))
- \*.dhc
- \*.dic (See [“Manual Migration of Dictionaries” on page 59.](#))
- \*.dir
- \*.eap
- \*.gen
- \*.ini
- \*.jsi
- \*.pro
- \*.rr
- radiusdata.\*

- **sbr\_id.xml**
- **\*.ses**
- **vista.\***

These subdirectories are also rolled forward:

- SBR JavaScripts in the **scripts** subdirectory

## Files from Earlier Releases That Require Manual Editing

Files listed in this section may require special handling during the migration process.

### Manual Migration of XML Configurations

You must manually move existing settings into the corresponding XML files that are delivered in the new Steel-Belted Radius Carrier software installation.



**CAUTION:** Do not modify any **\*.xml** files except these four:

- **/radius/sbr\_administration.xml**—Socket port number specified by port= setting
- **/radius/sbr\_ccm.xml**—Socket port numbers specified by port= settings
- **/radius/sbr\_id.xml**—Hostname specified by id= setting
- **radius/system/config/logging\_mgr.xml**—LogStream events= settings
- **radius/snmp/conf/events.xml**—Dilution factor specified by dilutionFactor= setting

The IP addresses, hostnames, socket port numbers, and other similar data entries found in these files often have corresponding parameters in the **radius.ini** file, which must be kept in agreement.

### Manual Migration of ROOT Certificates

The storage of root certificates is managed by the Steel-Belted Radius Carrier server. You use the Web GUI to add and delete root certificates. You must manually migrate root certificates, using the Web GUI to add them from the old root directory.

## Manual Migration of Dictionaries

If you have stored any modified or third-party dictionary files (\*.dci, \*.dcm, \*.dct, and \*.dic) in the /radius directory, you must manually migrate these either by merging each of the modifications with the corresponding files that are shipped with the new Steel-Belted Radius Carrier software, or by copying the third-party dictionary files to the new /radius directory.

If you are using the optional WiMAX mobility module, check the **vendor.ini** file to confirm that the WiMAX vendor ID is defined. If the field is blank, restore a file from the /install/backup/<time-stamp> directory.

## Manual Migration of Third-Party Plug-ins and Other Binaries

If you have stored any third-party plug-ins (\*.so) or other binaries in the old release's *radiusdir* directory, then you must manually migrate them by copying the files to the new *radiusdir*.

## Manual Migration of SNMP Files

We recommend that you not attempt to migrate a previous release's SNMP configuration, except if you have stored any modified or third party MIB files in the old server's **radius/snmp/mibs** directory. Migrate these files manually by copying them to the new server.

## Migrating from SBR Release 6.X

No files from SBR Release 6.X require special handling.

## Migrating from SBR Carrier Release 7.0

The only files that require special handling when migrating from SBR Carrier Release 7.0 to SBR Carrier Release are the files associated with storing CDR records when running in standalone mode.

In SBR Carrier Release 7.0, the settings for storing CDR records in a remote LDAP server were configured in the **ss7ldapdb.gen** file. When running SBR Carrier Release in standalone mode, these settings are configured in the **ss7db.gen** file. When migrating from Release 7.0 to Release, you need to move these settings from the **ss7ldapdb.gen** file to the **ss7db.gen** file.

## Migrating from SBR Carrier Release 7.4.0

SBR Carrier Release 7.4.0 uses the latest version of the RAIMA RDM Embedded Database. The RAIMA RDM Embedded database is upgraded from RDME version 7.0 to 11.0 when running the configure script.

**NOTE:** The RAIMA RDM Embedded database upgraded to the new version 11.0 cannot be downgraded to the old version. Make sure that the old version of the database is backed up. The following files must be backed up:

- radiusdata.dbd
- radiusdata.d01
- radiusdata.d02
- radiusdata.d03
- radiusdata.k02
- radiusdata.k01
- vista.taf

## Migration and New Installations of SBR Carrier with WiMAX

Previous releases of SBR Carrier introduced improvements in the WiMAX processing software. This change improves both performance and scalability. The improvements include different logic for assigning primary keys to WiMAX tables and for generating the Class attribute in the Access-Accept response. To control these improvements, a new parameter, EnableWiMAXUniqueSessionIdFromNAI, has been added to the **radius.ini** file.

**NOTE:** When the EnableWiMAXUniqueSessionIdFromNAI parameter is enabled, new session records in the database and the Class attribute in Access-Accept messages are incompatible with the WiMAX logic in previous releases of SBR Carrier. For this reason, this parameter is disabled by default.

For new installations of SBR Carrier software running the optional WiMAX module, we recommend that you enable this parameter as part of your installation procedure. It must be enabled on all SBR Carrier servers in your network.

If you are currently running a previous release of SBR Carrier and using WiMAX, the new WiMAX logic is backward compatible with the previously released WiMAX logic; that is, it understands the older session records and Class attributes. However, the older WiMAX logic does not comprehend the new session records or Class attribute in the Access-Accept messages. Due to this incompatibility, when you have multiple SBR Carrier servers in your network, you need to upgrade each server before enabling the EnableWiMAXUniqueSessionIdFromNAI parameter.

**NOTE:** If you are running the Session State Register (High Availability) cluster, this upgrade does not require any changes to the session database schema, and does not require any stopping or starting of nodes in the cluster. Therefore, it is nondisruptive.

For networks running multiple SBR Carrier servers, follow this procedure to upgrade:

**NOTE:** Upgrading the SBR Carrier software does not interrupt service.

1. Upgrade the first server with the latest SBR Carrier release software. Do not enable the EnableWiMAXUniqueSessionIdFromNAI parameter.
2. Upgrade the next server with the latest SBR Carrier release software. Do not enable the EnableWiMAXUniqueSessionIdFromNAI parameter.
3. Repeat until all servers are running the latest SBR Carrier release software.

At this point, all servers are running the latest SBR Carrier release software, but the EnableWiMAXUniqueSessionIdFromNAI parameter is disabled on all servers. The next step is to enable the EnableWiMAXUniqueSessionIdFromNAI parameter on all servers one by one.

1. Select the first server.
2. Edit the [Configuration] section of the **radius.ini** file and set EnableWiMAXUniqueSessionIdFromNAI=1.
3. Restart the server (sbrd restart). (The EnableWiMAXUniqueSessionIdFromNAI parameter is not updated on receipt of a HUP.)

After the server is back up, repeat this process on each and every server in the network one by one.

## Historic Files

When Steel-Belted Radius Carrier Release is installed and a previous release's files are migrated forward, the Release installation script creates another copy of some critical older files in the **/radius/install/backups**

directory. The data files in [Table 15 on page 62](#) record version and other operating information that is used by subsequent installations.



**CAUTION:** Do not move, rename, or otherwise modify these files.

**Table 15: Historic Files and Directories**

File	Function
<code>radius/install/package.dat</code>	Contains a unique package identifier.
<code>radius/install/preinstall.dat</code>	Contains the absolute path name of the backup directory for your old Steel-Belted Radius Carrier software and configuration files (pre-installation backup).
<code>radius/install/install.dat</code>	Contains the absolute path name of the backup directory for your Steel-Belted Radius Carrier software and configuration files, as shipped without modification (post-installation backup).
<code>radius/install/update.dat</code>	Contains the absolute place name of the Steel-Belted Radius Carrier migrate source, if any exist.
<code>radius/install/configure.dat</code>	Contains configuration state data.
<code>radius/install/uninstall.dat</code>	Contains the absolute place name of the backup directory for your Steel-Belted Radius Carrier software and working configuration files (pre-uninstall backup).
<code>radius/install/backups/</code>	Contains the backups referenced by the <code>.dat</code> files.



# 3

PART

## Installing and Configuring a SBR Carrier Standalone Server

---

Installation and Basic Configuration of a SBR Carrier Standalone Server | 64

---

# Installation and Basic Configuration of a SBR Carrier Standalone Server

IN THIS CHAPTER

- Unpacking the Steel-Belted Radius Carrier Software | 64
- Running the Steel-Belted Radius Carrier Configure Script | 68
- Starting and Stopping a Standalone Steel-Belted Radius Carrier Server | 77
- Basic SBR Carrier Configuration | 84
- Configuring the CST for a Standalone Server | 88

This chapter describes how to run the configure script that installs Steel-Belted Radius Carrier software on a standalone server, and the basic information about setting up a new Steel-Belted Radius Carrier server.

These topics are in the chapter:

## Unpacking the Steel-Belted Radius Carrier Software

You need to perform this task on each standalone server.

### Package Management Commands

The installer for the Solaris version of the Steel-Belted Radius Carrier server software uses **pkgadd** files, which have filenames that include the edition and version of the server software.

[Table 16 on page 64](#) lists common Solaris package management commands.

Table 16: Useful Package Management Commands

Command	Function
pkginfo -x  grep "JNPR"	Report any pre-existing packages and patches.

Table 16: Useful Package Management Commands (*continued*)

Command	Function
<code>pkginfo -l JNPRsbr</code>	Report high level description for specified package
<code>pkginfo -r JNPRsbr</code>	Show installed directory
<code>pkgadd -d /path/to/JNPRsbr.pkg [-a none] JNPRsbr.pkg</code>	Installs [with user prompts].  <b>NOTE:</b> You can install SBR at any location. We recommend you to install SBR at the default location <code>"/opt/JNPRsbr"</code> .
<code>pkgrm JNPRsbr</code>	Uninstall Steel-Belted Radius Carrier.

To install the Steel-Belted Radius Carrier server software:

On the Solaris server:

1. Log in to the Solaris server as root.
2. Copy the Steel-Belted Radius Carrier installation files from their download location to the Solaris server. Make sure to copy them to a local or remote hard disk partition that is readable by root.

This example copies the files from a download directory to the `/tmp/sbr` directory.

Execute:

```
mkdir -p /opt/tmp
cp -pR /tmp/sbr/solaris/* /opt/tmp/
```

3. Extract the Steel-Belted Radius Carrier installation package.

For 64-bit Solaris, execute:

```
cd /opt/tmp/
$ls -ltr
```

```
total 216240
-rw-r--r--  1 root      root      110712276 Aug 25 09:44
sbr-sa-8.6.0.R-1.sparcv9.tgz
```

Execute:

```
gunzip sbr-sa-8.6.0.R-1.sparcv9.tgz
tar xf sbr-sa-8.6.0.R-1.sparcv9.tar
```

4. Verify that the extraction worked and confirm the name of the package file.

For 64-bit Solaris, execute:

**ls -ltr**

```
total 216256
drwxr-xr-x  4 Xtreece  other      370 Aug 24 17:01 JNPRsbr.pkg
-rw-r--r--  1 root     root       110712276 Aug 25 09:44
sbr-sa-8.6.0.R-1.sparcv9.tar
```

5. Install the package.

Execute:

**pkgadd -d /opt/tmp** (location of the unzipped and untarred installation)

```
The following packages are available:
  1  JNPRsbr.pkg      JNPRsbr - Juniper Networks Steel-Belted Radius Carrier
                        (sparc) 8.60.50006  (this is an example only)
Select package(s) you wish to process (or 'all' to process all packages).
(default: all) [?,??,q]: all
```

6. Type **all** and press Enter.

The script resumes.

```
Processing package instance <JNPRsbr.pkg> from </tmp>
```

7. Confirm the working directory.

Depending on the system configuration, the script prompts you to create the **/opt/JNPRsbr** directory if it does not exist, to over-write an already extracted package, or any of several other questions.

```
The selected base directory </opt/JNPRsbr> must exist before installation is
attempted.
Do you want this directory created now [y,n,?,q]
```

8. Confirm that you want the directory created, if required. Enter **y**.

The script resumes.

```
Using </opt/JNPRsbr> as the package base directory.
#Processing package information.
#Processing system information.
#Verifying disk space requirements.
#Checking for conflicts with packages already installed.
```

```
#Checking for setuid/setgid programs.
This package contains scripts which will be executed with super-user
permission during the process of installing this package.
Do you want to continue with the installation of <JNPRsbr> [y,n,?]
```

9. Confirm that you want to continue to install the package. Enter **y**  
The script resumes.

```
Installing JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier) as <JNPRsbr>
#Executing preinstall script.
#Installing part 1 of 1.
/opt/JNPRsbr/radius/3GPP.dct
/opt/JNPRsbr/radius/3GPP2.dct
/opt/JNPRsbr/radius/3comsw.dct
/opt/JNPRsbr/radius/3gpp.ini
....
/opt/JNPRsbr/radius/wimaxAttributeProcessor.ctrl
/opt/JNPRsbr/radius/wimaxAttributeProcessor.so
/opt/JNPRsbr/radius/xylan.dct
[ verifying class <none> ]
#Executing postinstall script.
Newly installed server directory will be backed up as:
/opt/JNPRsbr/radius/install/backups/2009:08:25-09:47:28
Installation of <JNPRsbr> was successful.
```

On the Linux server:

1. Log in to the Linux server as root.
2. Copy the Steel-Belted Radius Carrier installation files from their download location to the Linux server.  
Make sure to copy them to a local or remote hard disk partition that is readable by root.

This example copies the files from a download directory to the **/tmp/sbr** directory.

Execute:

```
mkdir -p /opt/tmp
cp -pR /tmp/sbr/linux/* /opt/tmp/
```

3. Execute **cd /opt/tmp/**.
4. Install the package.

To install the 64-bit SBR package, execute:

**yum localinstall sbr-sa-8.6.0.R-1.el7.x86\_64.rpm** on RHEL 7.3 or later.

See “Linux” on page 47 for supported RHEL versions.

5. Installation proceeds and SBR is installed at `/opt/JNPRsbr/radius/install`.

## Running the Steel-Belted Radius Carrier Configure Script

1. Navigate to the directory where you installed the Steel-Belted Radius Carrier package (`/opt/JNPRsbr/radius/install`).

```
cd /opt/JNPRsbr/radius/install
```

2. Execute the **configure** script to install the Steel-Belted Radius Carrier server software:

```
./configure
```

3. Review and accept the Steel-Belted Radius Carrier license agreement.

Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

```
Do you accept the terms in the license agreement? [n]y
```

4. Enter the license key or keys that you collected in the preparatory step *Obtaining License Keys*. Then press Enter.

The **configure** script creates a license file that contains the key on the server for your Steel-Belted Radius Carrier software.

```
-----
SBR 8.60.50006
on SunOS 5.10 Generic_141444-09 node sbrha-10.spgma.juniper.net
-----
```

```
Enter SBR licenses meant only for this particular SBR node.
Enter one license per line and an empty line when finished.
Enter SBR full license: xxxx xxxx xxxx xxxx
Enter SBR feature license:
```

5. The script warns that you are changing the server configuration and prompts you to continue. Enter **y** to continue.

```
Generating configuration files
```

```
WARNING: You are about to make irreversible changes to this node.
```

```
Are you sure that you wish to continue? (y,n): y
```

## Migration

6. The script prompts for the type of installation, either a new installation or a migration from an earlier release:

```
Please enter backup or radius directory from which to migrate.
```

```
Enter n for new configuration, s to search, or q to quit
```

```
[n]: n
```

- For a new installation, enter **n**.
- If you are migrating an existing Steel-Belted Radius Carrier installation and have copied a previous release's files to the Release 8.6.0 server (in [“Creating a Copy of Existing SBR Server Release Files for Migration” on page 54](#)), enter the directory path to the copy of the old installation.

For example:

```
[ /opt/JNPRsbr/PreviousInstall ]:
```

- If you are migrating an existing Steel-Belted Radius Carrier installation and have copied a previous release's files to the Release 8.6.0 server (in [“Creating a Copy of Existing SBR Server Release Files for Migration” on page 54](#)), but you need to search for the directory that contains the Steel-Belted Radius Carrier files, enter **s**.

**NOTE:** Steel-Belted Radius Carrier Release 8.6.0 supports importing configuration files from previous versions of SBR Carrier.

If you select this option, remember that some files require manual editing and updating after installation. See [“Migrating from Previous SBR Releases” on page 56](#).

7. Specify whether you want to use the JRE installed in your system to enable JDBC plug-ins and JavaScript implementation.

```
Do you want to configure Java Runtime Environment for JDBC Feature [n] :
```

- If no, press Enter to proceed to the next prompt. SBR Carrier does not support JDBC plug-ins unless you specify a valid JRE path.
- If yes, type **y** and press Enter. You are prompted to specify the path where the JRE is installed in your system. The Java Virtual Machine (JVM) architecture should be compatible with SBR Carrier.

**NOTE:** Java 1.8.0 or a later version is required to access the Web GUI. To support both JDBC plug-ins and Web GUI, it is recommended to use Java 1.8.0 or a later version with the JVM architecture compatible with your SBR Carrier. For example, if you are using the 64-bit version of SBR Carrier, you must use the 64-bit version of Java 1.8.0 or later.

```
Enter 64-bit libjvm.so path (Ex: /opt/jvm/jre/lib/amd64/server/ ) :
```

**NOTE:** If you enter an incorrect JVM path three times, SBR Carrier proceeds to the next step. In this case, you will not be able to use JDBC plug-ins. To specify the valid JVM path, you need to run the configure script again.

8. Supply the name of the initial admin user, root.

```
Enter initial admin user (UNIX account must have a valid password) [root]:
```

Press Enter to accept the default, root.

9. Specify whether you want to set up centralized configuration management (CCM).

```
Enable Centralized Configuration Management (CCM) for this SBR node? [n]: y
```

10. If you have enabled CCM, specify the server role.

```
Configure SBR node as CCM primary (p) or replica (r)? [r]: p
```

- a. If no other servers have been installed with CCM enabled, this is the first server to be installed, and CCM is used, this server must be the primary. Enter **p** to set the role as primary.



If you set the role to primary, the script prompts for the host secret string.

**Enter primary host secret:**

Type the secret string and press Enter.

- b. The script prompts you to confirm the host secret string

**Confirm primary host secret:**

Type the secret string again and press Enter.

11. Specify whether you want to use the auto-restart module that automatically restarts the SBR Carrier server in case of an unexpected shutdown.

```
Do you want to enable "Radius WatchDog" Process? [n]: Y
Radius WatchDog feature set to Enable
Please ensure that Perl 5 or better is installed.
```

**NOTE:** If Perl version 5 is not installed, the radiusd script will not run, even if enabled by configuration, and SBR Carrier will operate without the auto-restart module running.

12. Specify whether you want to configure SBR Carrier to provide LDAP server emulation for configuration and statistics using the LCI.

```
Do you want to enable LCI? [n] :
```

- If no, press Enter to accept the default.
- If yes, enter **y** and press Enter. You are prompted to provide information for LCI configuration.
  - a. When you are prompted for the port number, enter the port number that is used for communication between SBR Carrier and the LDAP client.

**NOTE:** SBR Carrier uses port 667 as the default for LDAP emulation to avoid conflict with other LDAP servers.

```
Configure LCI Port [667]: 1026
```

- b. The script displays the interfaces available in the system. When you are prompted to enter interface addresses on which Steel-Belted Radius Carrier should listen for LCI requests, enter the addresses you want to use from the Available Interfaces list.

```

LCI Interface Configuration :
Available interfaces :
127.0.0.1
10.212.10.66

HELP : Enter one interface per line and an empty line when finished.

Enter LCI interface addresses from the above list.

Enter LCI interface address : 10.212.10.66
Enter LCI interface address : 127.0.0.1
Enter LCI interface address :

```

**NOTE:** SBR Carrier uses all interfaces for listening to LCI requests if you do not enter any interfaces.

- c. Specify whether you want to change the default LCI password to prevent unauthorized LDAP clients from accessing your database.

```

Do you want to change LCI Password? [n]:

```

- If no, press Enter to accept the default password.
- If yes, enter y and press Enter. You are prompted to enter a new password.

```

Do you want to change LCI Password? [n]: Y

Password must meet the following requirements:

1. 6-8 Alphanumeric characters.
2. No Special characters other than underscore ('_').

Enter Password:
Confirm Password:
Password will be changed when SBR restarts.

```

**NOTE:** Make sure that the entered password is at least 6 alphanumeric characters and not more than 8 characters in length. The password should not include any special characters other than underscore ('\_').

**NOTE:** The configure script also checks whether the LDAP utilities (such as **ldapdelete**, **ldapmodify**, and **ldapsearch**) are installed in your system. For Linux, a warning message is displayed if you have not installed any of these utilities in your system. For Solaris, LDAP utilities are shipped with SBR Carrier package.

13. Specify whether you want to configure Steel-Belted Radius Carrier for use with an Oracle database.

To support this option, the server must already be configured as an Oracle client. (See [“Setting Up External Database Connectivity \(Optional\)” on page 53.](#))

```
Configuring for use with generic database
Do you want to configure for use with Oracle? [n]:
```

If no, press Enter to accept the default.

If yes, type **y** and press Enter. You are prompted for version and path information for the Oracle library files.

```
Do you want to configure for use with Oracle? [n]: y
Supported Oracle versions: 10, 11, 12
What version of Oracle will be used? [10]: 10
Configuring for use with Oracle 10
Setting the environment variable ORACLE_HOME
Enter ORACLE_HOME [/dbms/u10/app/oracle/product/10.2.0]:
Setting the environment variable LD_LIBRARY_PATH
Enter path for Oracle shared libraries
[/dbms/u10/app/oracle/product/10.2.0/lib]:
Setting the environment variable TNS_ADMIN
Enter TNS_ADMIN [/dbms/u10/app/oracle/product/10.2.0/network/admin]:
```

**NOTE:** You must configure 64-bit Oracle client for 64-bit SBR Carrier.

14. Specify whether you want the Steel-Belted Radius Carrier server to communicate with an SS7 system using SIGTRAN.

To support this option, the server must already be configured to support SIGTRAN using Signalware. (See [“Installing the SIGTRAN Interface \(Optional\)” on page 55](#) for an overview, and [“SIGTRAN Support for Steel-Belted Radius Carrier” on page 309](#) for specific instructions.)

```
Do you want to configure for use with SIGTRAN? [n]: y
Configuring for use with SIGTRAN
Setting the environment variable OMNI_HOME
Enter OMNI_HOME [/opt/JNPRss7]:
```

15. Specify whether you want to start the GWrelay process while executing the `./sbrd start` script.

```
Do you want to enable "GWrelay" Process? [n]: y
GWrelay will be started with sbrd
```

16. Specify whether you want to install the optional SNMP module to monitor your Steel-Belted Radius Carrier server from an SNMP management station.

```
Do you want to configure SNMP? [n]:
```

- If no, press Enter to proceed to the next prompt.
- If yes, type **y** and press Enter. The installer prompts you for the information it needs to configure the `jnrpsnmpd.conf` and `startsnmp.sh` files.
  - a. When you are prompted for a community string, enter the community string used to validate information sent from the SNMP subagent on the Steel-Belted Radius Carrier server to your SNMP management station.  
Choose a community string: **public**
  - b. When you are prompted for a range of IPv4 addresses, specify a starting IP address in Classless Inter-Domain Routing format. To specify that only one host may query the agent, enter the IP address of the host followed by `/32`. To specify that any host on a designated class C network may query the agent, enter the starting address of the network followed by `/24`.

```
Specify the range of IPv4 addresses that may query this agent, such as
1.2.3.0/24.
Address range: 192.168.70.0/24
```

- c. If you are using SNMPv2, enter the DNS name or IP address of the trap sink to receive trap information from the SNMP subagent on the Steel-Belted Radius Carrier server.

```
SNMPv2 trap sink: 192.168.70.86
```

- d. Set the SNMP agent port.

Although you may specify the default SNMP port, 161, we recommend that you specify a different port to avoid contention with other agents that are likely to already be using 161. If you choose an alternate port, make a note of it because your MIB browser needs to be configured to the same setting.

```
Specify SNMP agent listening port[161]: 24161
```

- e. Specify a trap sink address, if required.

```
Optionally specify a trap sink address that will receive SNMPv2 trap
[localhost]: 172.28.72.83 2
SNMPv2 trap sink port[162]:
Configuration of SNMP complete.
```

17. The script searches for the Java 1.8.0 or later version in the default system path and displays a confirmation message if found.

```
Configuring Admin GUI Webserver
Compatible Java version 1.8.0_66 found in: /usr/java/jdk1.8.0_66
```

If the specific version is not found, the script prompts you to enter the directory path where the specific Java version is installed in your system. For example, if the Java version is available in the **/usr/java/jdk1.8.0\_66/bin/java** path, enter only **/usr/java/jdk1.8.0\_66**, do not include **bin/java**.

```
Enter Java version 1.8 installed path :
```

18. Specify whether you want to install a custom SSL certificate for the Web GUI. For more information about certificates, see the *SBR Carrier Administration and Configuration Guide*.

```
Do you want to install custom SSL certificate for Admin WebServer? [n]:
```

- If no, press Enter. A self-signed certificate is created and installed in the web server.
- If yes, enter **y** and press Enter. You are prompted to enter the absolute path where the SSL certificate is available. For example, **/opt/customSSLCert.pfx**.

```
Enter the absolute path to certificate.
Note: Only *.pfx files are accepted. (Example-/opt/customSSLCert.pfx):
```

When you are prompted for the password, enter the password to open the SSL certificate.

```
Enter the password to open the certificate :
```

13. Specify whether you want to configure the Steel-Belted Radius Carrier server to autoboot (restart automatically when the operating system is restarted).

```
Enable (e), disable (d), or preserve (p) RADIUS autoboot [e]: e
```

Steel-Belted Radius Carrier stores its autoboot settings in the local **/radiusdir/radius/sbrd** file.

- If you enter **e** (enable), the **configure** script copies the local **sbrd** script to **/etc/init.d**, where it is automatically invoked by the OS whenever the OS is stopped or started.
- If you enter **d** (disable), the **configure** script removes all copies of the **sbrd** script from **/etc/init.d**, thus, disabling autoboot for all versions of Steel-Belted Radius Carrier.
- If you enter **p** (preserve), the **configure** script does nothing, thereby leaving your previous autoboot scripts unchanged.

When you finish entering settings, the script configures Steel-Belted Radius Carrier with the specified settings and then displays:

```
The SBR Admin Web GUI can be launched using the following URL:
https://<servername>:2909

Configuration complete
```

## Starting and Stopping a Standalone Steel-Belted Radius Carrier Server

After you have successfully run the configure script, you need to start the Steel-Belted Radius Carrier server. To begin, start the SBR daemon using `sbrd` (see “[sbrd](#)” on page 77). Make sure that you can stop it and check the server status from a terminal command line.

### `sbrd`

You use the `sbrd` script to start, stop, or restart the RADIUS process (and LDAP process only if SS7 is configured) on standalone Steel-Belted Radius Carrier servers. The LDAP process (`slapd`) is used to store session information only if SS7 is configured. The `sbrd` script may be in either of two directories on servers, depending on whether they have been configured to automatically start all procedures or not using the autoboot functionality that is configured when running the `configure` script. All `sbrd` commands are executed by root.

#### *Running `sbrd` on a Standalone Server*

##### **Syntax**

Executing start, stop, or restart always starts the RADIUS process. Whether the LDAP process starts when you execute start, stop, or restart depends on whether you enable or disable SS7 while running the configure script. If you answer `n` when prompted “Do you want to configure for use with SIGTRAN? [n]:”, SS7 is disabled and the LDAP process is not started when you execute start, stop, or restart. In this case, the available `sbrd` usage is:

```
sbrd status
sbrd start      [force]
sbrd stop       [force]
sbrd restart    [force]
sbrd clean      [force]
sbrd hup
sbrd status -v [-p <LCI password>]
```

However, if you answer `y` when prompted “Do you want to configure for use with SIGTRAN? [n]:”, SS7 is enabled and the LDAP process is started when you execute start, stop, or restart. In this case, the available `sbrd` usage is:

```
sbrd status      [radius|ss7ldapdb]
sbrd start        [radius|ss7ldapdb] [force]
sbrd stop         [radius|ss7ldapdb] [force]
sbrd restart      [radius|ss7ldapdb] [force]
sbrd clean        [radius|ss7ldapdb] [force]
```

```
sbrd hup      [radius|ss7ldapdb|authGateway [process-name]]
sbrd status  [radius|ss7ldapdb] -v [-p <LCI password>]
```

If SS7 is enabled, the **start**, **stop**, and **restart** arguments start, stop, and stop and restart both the RADIUS and LDAP (slapd) processes on the local Steel-Belted Radius Carrier server. For example, invoking **sbrd start** starts both the RADIUS and LDAP (slapd) processes on the local SBRC server. You can also use these arguments with either the **radius** or **ss7ldapdb** option to individually start, stop, and stop and restart the RADIUS and LDAP (slapd) processes on the local SBRC server. For example **sbrd start radius** starts just the RADIUS process on the local server.

### Options

- The **clean** argument removes lock files that prevent reinitializing the database more than once. You should use this argument only if things go wrong during the initial installation and configuration.
- The **hup** option operates as the **kill -HUP** command does on SBR Carrier nodes, but does not require the process ID. Executing **sbrd hup authGateway** issues the SIGHUP (1) signal to all the authGateway processes running on SBR Carrier. To issue the SIGHUP (1) signal only to the specific authGateway process, you must execute the **hup** option with the authGateway process name, for example: **sbrd hup authGateway GMT**.
- The **force** argument makes **sbrd** attempt to disregard or overcome any errors that occur when processing the command. Normal behavior without the argument is to halt on errors. For example, **sbrd start** does not attempt to start software that is already running, but **sbrd start force** ignores a running process. This may produce unintended results, so use **force** with great care.
- The **-v** option displays additional information about the RADIUS process along with basic information such as the SBR package version, SBR process status, and SBR process ID. If you have changed the default Lightweight Directory Access Protocol (LDAP) Configuration Interface (LCI) password, you should use the **-p** option to specify the password. For more information about the RADIUS status information, see [“Displaying RADIUS Status Information” on page 79](#).

## Starting the RADIUS Server

To start the RADIUS and LDAP processes manually, execute as root:

```
cd radiusdir
./sbrd start
```

If you change configuration settings for your Steel-Belted Radius Carrier server, you may need to restart Steel-Belted Radius Carrier to make the changes effective. As an alternative to issuing a **sbrd stop** command immediately followed by a **sbrd start** command, you can use the **sbrd restart** command to restart Steel-Belted Radius Carrier. When you issue the **sbrd restart** command, Steel-Belted Radius Carrier shuts down and then immediately restarts the RADIUS processes.



```
cd radiusdir
./sbrd restart
```

## Stopping the RADIUS Server

Use the following commands to stop the RADIUS server:

```
cd radiusdir
./sbrd stop
```

When you execute the `sbrd stop` command, Steel-Belted Radius Carrier allows its subsystems to complete outstanding work and release resources, and then stops the RADIUS processes gracefully.

If Steel-Belted Radius Carrier fails to stop after you issue the `sbrd stop` command, you can use the optional `force` argument to terminate all subsystems immediately.

```
cd radiusdir
./sbrd stop force
```

## Displaying RADIUS Status Information

You can use the following command to display basic information (such as SBR package version, SBR process status, and SBR process ID) about the RADIUS process:

```
cd radiusdir
./sbrd status
```

The system responds with:

```
-----
-----
SBR-64 8.60-R1.0
on SunOS 5.10 Generic_141444-09 node uranus.carrier.spgma.juniper.net
-----
172.28.84.73.1646 Idle
172.28.84.73.1813 Idle
172.28.84.73.1645 Idle
172.28.84.73.1812 Idle
*.1813 *.* 0 0 49152 0 LISTEN
*.1812 *.* 0 0 49152 0 LISTEN
```

```

root 6628 ./slapd -h ldap://127.0.0.1:389 -f
/opt/JNPRsbr/radius/openldap/slapd.conf
root 4449 radius sbr.xml
root 1189 webserver

```

You can use the **sbrd status** command with the **-v** option to display the following additional information about the RADIUS process along with the preceding information:

- **Loaded Plug-in Information**—Displays the name, version, and status of the loaded authentication and accounting plug-ins. The InitializationString value of the plug-in is displayed as the name of the plug-in.
- **License Status**—Displays the license key, feature name, and license status with expiry date.
- **IP Pool Information**—Displays the pool name, IP range of the pool, total number of addresses in the pool, and total number of available addresses in the pool.

**NOTE:** This information is displayed only for the SBR standalone version.

- **IP Ranges and IP Caches**—For IP ranges, displays the pool name, start address of the pool, end address of the pool, and total number of addresses in the pool. For IP caches, displays the pool name, total number of addresses in the pool, and percentage of the addresses in the pool that are cached.

**NOTE:** This information is displayed only for the SBR cluster version.

- **Statistics Information**—Displays statistical information about SBR such as current sessions count, SBR uptime (in seconds), current rate details, and average transaction rate since the SBR server started. The transaction rate is calculated using the following formula:

Transaction Rate = Total Transaction Count / SBR Running Time

where:

Total Transaction Count = Total Authentication Transactions + Total Accounting Transactions

**NOTE:** The current rates are updated at a time interval of one second.

- **Radius Ports Information**—Displays protocols, port types, and port numbers on which the SBR Carrier server is listening.
- **Proxy Configuration Information**—Displays proxy names and their IP addresses.

- CST Store Information—Displays the active session store and the time when the last session persistence switchover occurred.
- authGateway Process Information—Displays all the active authGateway process information.

**NOTE:** To display this additional information, you need to enable the Lightweight Directory Access Protocol (LDAP) Configuration Interface (LCI) in the **radius.ini** file. If you have changed the default LCI password (which we strongly recommend), you should use the **-p** option to specify the password, that is **./sbrd status -v -p [LCI password]**.

To display additional information about the RADIUS process along with the basic information, execute the following command:

```
cd radiusdir
./sbrd status -v
```

The system responds with:

```
-----
-----
SBR-64 8.60-R1.0
on SunOS 5.10 Generic_141444-09 node xyz.juniper.net
-----
172.28.84.73.1646 Idle
172.28.84.73.1813 Idle
172.28.84.73.1645 Idle
172.28.84.73.1812 Idle
*.1813 *.* 0 0 49152 0 LISTEN
*.1812 *.* 0 0 49152 0 LISTEN

root 6628 ./slapd -h ldap://127.0.0.1:389 -f
/opt/JNPRsbr/radius/openldap/slapd.conf
root 4449 radius sbr.xml
root 58284 GWrelay-64
root 1189 webserver
root 21316 /opt/JNPRsbr/radius/authGateway -name GMT -conf
/opt/JNPRsbr/radius/conf/authGateway.conf -start
root 21317 /opt/JNPRsbr/radius/authGateway -name GMT1 -conf
/opt/JNPRsbr/radius/conf/authGateway.conf -start
root 21318 /opt/JNPRsbr/radius/authGateway -name GMT2 -conf
/opt/JNPRsbr/radius/conf/authGateway.conf -start
Radius Ports Information :
```

Protocol	Port Number	Port Type
TCP	1812	Admin
TCP	64277	Others
TCP	64278	Others
TCP	5235	Others
TCP	64281	Others
TCP	64229	Others
UDP	28000	Proxy
UDP	28001	Proxy
UDP	28002	Proxy
UDP	28003	Proxy
UDP	28004	Proxy
UDP	28005	Proxy
UDP	28006	Proxy
UDP	28007	Proxy
TCP	1813	Admin
TCP	1814	Others
UDP	1646	Radius
UDP	1813	Radius
UDP	1645	Radius
UDP	1812	Radius

#### IP Pools Information :

Pool Name	Ip Address Range	Total	Available
POOL1	10.10.10.1:20	20	20
POOL2	20.20.20.1:20	20	20

#### Statistics Information :

```

Current Sessions Count           = 94831
Transaction Rate                 = 62 TPS (Since Server Start)
SBR Uptime                       = Up Since 2013/06/19 06:39:19 [ 1 Hrs
    : 59 Mins : 56 Secs ]
auth-request-current-rate        = 965
auth-accept-current-rate         = 926
auth-reject-current-rate         = 914
acct-start-current-rate           = 0
acct-interim-current-rate         = 0

```

```

acct-stop-current-rate           = 0
proxy-auth-request-current-rate  = 0
proxy-acct-request-current-rate  = 0
proxy-fail-timeout-current-rate  = 0
proxy-fail-badresp-current-rate  = 0
proxy-fail-badsecret-current-rate = 0
proxy-fail-missingresr-current-rate = 0
proxy-retries-current-rate       = 0
proxy-auth-rej-proxy-current-rate = 653
proxy-acct-fail-proxy-current-rate = 685
proxy-auth-rej-proxy-error-current-rate = 0
proxy-transaction-current-rate   = 0

```

#### Proxy Configurations :

Proxy Name	Ip Address
KIX	10.13.20.62
TRIX	10.13.20.60

#### Loaded Plugins Information :

Plugin Name	Plugin Version	Status
LDAP	v7.5.0.A-0.0	Success
SQL-ORACLE	v7.6.0.B-0.0	Success
SIMAUTH	v7.6.0.B-0.0	Success
EAP-TLS	v7.6.0.B-0.0	Success
EAP-TTLS	v7.6.0.B-0.0	Success
CDRACCT	-	Failed
SQL-ORACLE-ACCT	v7.6.0.B-0.0	Success
SQL-JDBC-ACCT	v7.6.0.B-0.0	Success
ldapaccessor	v7.6.0.B-0.0	Success
sqlaccessor	v0.0.0.B-0.0	Success

#### CST Store Information :

Active Store : NDB

Last swapped @ Mon Sep 22 18:46:12 2014

#### License Information :

Expiry Date	License Key	Feature Name	Status
12/31/2015	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	SIM Authentication feature	Valid
-	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	Run license	Maint_Upgrade - Valid
12/31/2024	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	SIM Authentication feature	Trial License - Valid
10/31/2010	xxxxxxxxxxxxxxxxxxxxxxxxxxxx	WiMAX Mobility feature	Trial License - Expired

## Basic SBR Carrier Configuration

### Launching Web GUI

To test whether the server can be accessed by a management workstation, launch the Web GUI.

To launch the Web GUI:

1. Open a browser connection to the SBR Carrier server you want to administer.
  - To administer a SBR Carrier server running on your local host, enter **https://localhost:2909**, where the port assignment of 2909 is the application's default TCP port for administration connections.
  - To administer a SBR Carrier server running on a remote host, enter **https://server:2909**, where the port assignment of 2909 is the application's default TCP port for administration connections. For example:

https://192.168.24.15:2909/

**NOTE:** Make sure that you access the Web GUI using HTTPS, instead HTTP.

The **Login** page ([Figure 12 on page 85](#)) appears.

Figure 12: Login Page

Steel-Belted Radius Carrier

JUNIPER NETWORKS

Login

Username:

Password:

Login Reset

**NOTE:** Consult Juniper Networks Technical Support for changing the port number. Using a non-default port may cause communication problems between Web GUI and the SBR Carrier server.

2. Enter your administrator username in the **Username** field.
3. Enter your login password in the **Password** field.
4. Click **Login**.

When you click **Login**, Web GUI establishes an HTTPS connection with the local or remote server. The Web GUI displays an error message if the connection cannot be established.

**NOTE:** If a timeout occurs, verify that the SBR Carrier daemon is running on the target server and that it is listening on the administration port you entered in the URL; that the port is not blocked.

Web GUI verifies that the username you entered exists in the **access.ini** file. If the username is found, Web GUI validates the password you entered against a local or remote password database.

When you connect to a server, the **Home** page lists various features of the running server, such as version, platform on which it is running, IP address, available authentication methods, license information, and any initialization errors that might have occurred.

## Configuring the Server

After you have installed the Steel-Belted Radius Carrier software on the server, have added the appropriate licenses, and can work with Web GUI, you can begin configuring the software. The specific steps you must perform depend on your network's authentication and accounting needs.

The basic steps for configuring the Steel-Belted Radius Carrier environment include:

1. Configure each of your RADIUS client devices to communicate with your Steel-Belted Radius Carrier server. To do this, you must log in to each device and run its configuration interface.
2. Use the **RADIUS Clients List** page in Web GUI to configure the server to communicate with each RADIUS client. Details are in the *Administering RADIUS Clients and Client Groups* section of the *SBR Carrier Administration and Configuration Guide*.
3. If the clients use RADIUS Location Groups or IP Address Pools, each of those entities must exist before the clients are configured. See the sections on *Administering RADIUS Location Groups* and *Administering Address Pools* in the *SBR Carrier Administration and Configuration Guide* for instructions on setting these up.
4. Use the Users panel to identify the users or groups of users who are permitted to access the RADIUS clients.

Specify user attributes by selecting them in the Users panel or by creating user profiles in the Profiles dialog.

For more information, see the *SBR Carrier Administration and Configuration Guide*.

## Configuring SNMP

Steel-Belted Radius Carrier runs its own SNMP agent, but other SNMP agents run on most servers. In general, only one application can use a socket port; they are not shared resources.



During installation, the **configure** script prompts you for SNMP setup information, including an opportunity to specify a port other than the default 161 port that is usually in use by the Solaris SNMP agent.

- If you already provided an alternate port during that setup step, you can skip the following procedure about how to change the port number, but remember to set your MIB browser to listen on the port you specified.
- If you know that other agents already use port 161 but you did not specify an alternate during installation, change the Steel-Belted Radius Carrier port assignment by editing both **radiusdir/snmp/conf/jnprsnmpd.conf** and **radiusdir/snmp/bin/testagent.sh**. Remember to check your MIB browser to determine whether it also needs adjustment to communicate with the SBR Carrier server.

To change the port, edit the SBR Carrier SNMP configuration files listed in [Table 17 on page 87](#):

Table 17: SNMP Configuration Files

Filename	Function
<b>jnprsnmpd.conf</b>	Stores settings for the Steel-Belted Radius Carrier SNMP agent.
<b>testagent.sh</b>	Test script that verifies the Steel-Belted Radius Carrier SNMP agent is operating correctly.

- Edit **radiusdir/snmp/conf/jnprsnmpd.conf** to change the port number.

The **jnprsnmpd.conf** file is self-documenting. For more information about this file, see the *SBR Carrier Reference Guide*.



**CAUTION:** The **jnprsnmpd.conf** file is very sensitive to stray white space and the order in which sections and parameters appear. Mistakes in this file can disable SNMP.

- Make sure to make a backup copy of the file before making any changes.
- While editing the file, do not to make any unnecessary changes. Follow the embedded examples as closely as possible.
- When specifying networks, as in 172.28.68.0/24 in the **com2sec mynetwork 172.28.68.0/24 public** line, the trailing 32-x bits of the IP address must be zero as specified by the trailing /x notation. For example, 32-24=8 bits in this case.

2. Make the same port number change in **radiusdir/snmp/bin/testagent.sh** script, which is used to test the agent.
3. After making the change, restart either the Steel-Belted Radius Carrier server process or just its SNMP daemon.

Execute:

`/etc/init.d/init.jnprsnmpd start`

4. If necessary, set up your SNMP browser to listen on the new port.
5. To verify that the jnprsnmpd SNMP agent functions, run the **`radiusdir/snmp/bin/testagent.sh`** script.

**NOTE:** Refer to the *SBR Carrier Administration and Configuration Guide* for more information about configuring the SBR SNMP agent.

## Configuring the CST for a Standalone Server

For a standalone server, the CST is in local memory, and is configured with the **`dbclusterlocal.gen`** file when you run the configuration script. The **`dbclusterlocal.gen`** file is a small file consisting of a [Bootstrap] section that includes `LibraryName = dbclusterlocal` and `Enable=1`.

In addition, the setting of the **`PersistSessions`** parameter in the **`radius.ini`** file determines whether sessions are restored or not restored when SBR Carrier is restarted. The default and recommended settings for a standalone server are:

```
LibraryName = dbclusterlocal (set by default)
Enable =1 in dbclusterlocal.gen (set by default)
Enable=0 in the dbclusterndb.gen file (set by default)
PersistSessions=local in the radius.ini file
```

This configuration enables the CST in local memory and saves session information in a local file on SBRC so that the information is available after the server restarts.

You cannot configure field names in the local CST (**`dbclusterlocal.gen`**). However, there are three predefined fields and seven generic fields you can configure using the **`sessionTable.ini`** file. The three predefined fields are:

- `WimaxSessionId`
- `AcctMultiSessionId`
- `FunkOuterUserName`

The seven generic fields are:

- `Generic1`
- `Generic2`
- `Generic3`

- Generic4
- Generic5
- Generic6
- Generic7

In order to use any of the ten predefined fields, they must be mapped to attributes by editing the **sessionTable.ini** file.

### Configuring the sessionTable.ini file

The **sessionTable.ini** file allows you to take any attribute in a request and store it in the CST. To configure the seven generic fields in the standalone CST, you use RadAttr fields. RadAttr fields are designed to capture data from RADIUS packet attributes. You must be knowledgeable about RADIUS attributes, how they are used in the network, and which attributes to incorporate into the CST.

Valid **Radius-Attribute-Names** occur in Steel-Belted Radius Carrier dictionary files, as embodied in the SBR Carrier internal primary dictionary (which includes attributes defined in files having **.dct**, **.dcm**, **.dci**, and **.dcx** extensions in the SBR Carrier installation root directory). For information about SBR Carrier dictionaries, refer to the **readme.dct** and **dictiona.dcm** files.

#### **RadAttr Fields and the sessionTable.ini File**

RADIUS attributes can be generated by Steel-Belted Radius Carrier at four *capture points*, events that apply only to RadAttr fields:

- [AuthRequest]
- [AuthResponse]
- [AcctRequest]
- [AcctResponse]

Each RadAttr CST field requires at least one capture point to identify when RADIUS attributes are loaded in the field. Specify one or more capture-points by placing the field entry in one or more of the four capture point sections of the **sessionTable.ini** file.

One **Radius-Attribute-Name** can be captured multiple times, at different capture-points, by using multiple (different) **RadAttrFieldNames**, all of which are mapped to the same **Radius-Attribute-Name**.

If a RADIUS attribute is not present when a session's RADIUS packet is scanned at a capture point event, the corresponding CST field is not updated. If the specified attribute does not occur in a packet at any capture point, the session displays a null value in the CST field.

If a session experiences multiple occurrences of packets at the same capture-point carrying the same **Radius-Attribute-Name**, the values of subsequent occurrences overwrite the values of preceding occurrences

in the **RadAttrFieldName** field. For example, if a session receives multiple Accounting-Start packets, then the AcctRequest and AcctResponse capture-points are exercised more than once for that session, with the result that the captured **Radius-Attribute-Name/ RadAddrFieldnames** are updated multiple times. Thus, the CST always reflects the most recent values from each capture-point.

### Creating RadAttr Field-to-Attribute Mappings

The **sessionTable.ini** file maps **RadAttrFieldName** RadAttr CST fields to **Radius-Attribute-Name** RADIUS attributes:

**RadAttrFieldName = Radius-Attribute-Name**

### Defining RadAttr Fields

To define a RadAttr field, edit the **sessionTable.ini** file.

Add the **RadAttrFieldName = Radius-Attribute-Name** field-to-attribute mapping line to the **sessionTable.ini** file in the appropriate capture point section or in several sections.

For example:

```
[AcctRequest]
;=====
```

```
Generic1=3GPP-IMSI
Generic2=3GPP-GGSN-Address
```

### Multi-Valued Attributes

The **sessionTable.ini** file supports a special syntax to support multi-valued attributes. The RADIUS protocol permits certain attributes to be *multi-valued*, meaning that two or more instances of the attribute can occur in a single RADIUS packet. System fields in the CST always involve single-valued attributes. RadAttr fields in the CST can be single-valued or multi-valued.

For RadAttr fields in the CST, we support multi-valued attributes with the six MVA facilities defined in the list that follows. Any field-to-attribute mapping (in **sessionTable.ini**) can have its **Radius-Attribute-Name** optionally appended with one (but no more than one) multi-valued attribute @ extension.

### MVA-Count

The MVA-Count extension stores the count (0, 1, or more) of attributes of the specified **Radius-Attribute-Name** (occurring in the RADIUS packet being scanned at the specified capture-point).

Syntax:

*RadAttrFieldName = Radius-Attribute-Name@#*

### **MVA-Index**

The MVA-Index extension stores the value of the Nth attribute (where N is equal to or greater than 1) of the **Radius-Attribute-Name** occurring in the RADIUS packet being scanned at the specified capture-point. For example, @1 means to store the first value of the specified **Radius-Attribute-Name**. If the packet carries fewer than N attributes of the designated type, a value of null is stored.

In this case, N is a 1-based, not 0-based, integer. Thus, N denotes a positive number in the range 1–2038, expressed in standard base-10 notation. Specifying a value for N outside this range is an error.

Syntax:

*RadAttrFieldName = Radius-Attribute-Name@N*

### **MVA-First**

The MVA-First extension stores the first value of the **Radius-Attribute-Name** attribute occurring in the RADIUS packet being scanned at the specified capture-point. MVA-First extension is synonymous with the **Radius-Attribute-Name@1** extension.

Syntax:

*RadAttrFieldName = Radius-Attribute-Name@^*

### **MVA-Last**

The MVA-Last extension stores the last value of the **Radius-Attribute-Name** attribute occurring in the RADIUS packet being scanned at the specified capture-point.

Syntax:

*RadAttrFieldName = Radius-Attribute-Name@\$*

### **MVA-CharString-Concat**

The MVA-CharString-Concat extension treats the potentially multiple values of **Radius-Attribute-Name** as text strings. It concatenates them using the specified string "...", as a field delimiter between successive values.

The value of the delimiter string cannot contain the 2-character sequence "@" (at-sign/double-quote) or internal NULL characters. The delimiter string can consist of any arbitrary (single-byte) characters, such as comma, colon, space, tab, and new line characters. This facility supports C/C++-language backslash-escaping (meta character) conventions. For example, to include a backslash character in a delimiter string or in a character data-value, enter it as: \\.

Syntax:

```
RadAttrFieldName = Radius-Attribute-Name@"..."
```

### **MVA-ByteString-Concat**

The MVA-ByteString-Concat extension treats the potentially multiple values of ***Radius-Attribute-Name*** as byte strings. It concatenates them using a length and value representation: every value is represented by a one-byte length header that identifies the length of the data value in the range 1–253 and is followed by the data value. The concatenation terminates with a one-byte 0-value trailer that indicates a data of length 0, such as the end of the concatenated byte-string.

Syntax:

```
RadAttrFieldName = Radius-Attribute-Name@*
```

### **Using MVA Facilities**

The syntax of the MVA facilities supports single-valued attributes and multi-valued attributes. For single-valued attributes, **@1**, **@^**, **@\$**, and no-**@**-extension returns the value of the attribute. The **@"..."** and **@\*** extensions return one item in the concatenation; no delimiter appears for **@"..."**, while the 0-byte terminator immediately follows the first item in the list for **@\***.

Any field-to-attribute mapping in **sessionTable.ini** (whether its ***Radius-Attribute-Name*** is single-valued or multi-valued) that does not contain an MVA **@** extension is considered to be equivalent to the same field-to-attribute mapping with an **@^** extension appended to it.

In both the **@"..."** and the **@\*** cases, the individual values in the MVA are listed in the order they occur in the RADIUS packet. According to the RADIUS protocol, attribute sequence is significant.

# 4

PART

## Installing Session State Register Nodes

---

Installing Session State Register Nodes | **94**

Basic SBR Carrier Node Configuration | **138**

Customizing the SSR Database Current Sessions Table | **144**

Other Session State Register Configuration Files | **182**

When and How to Restart Session State Register Nodes, Hosts, and Clusters | **196**

---

# Installing Session State Register Nodes

## IN THIS CHAPTER

- Cluster Installation Overview | 94
- Unpacking Session State Register Software | 95
- Setting Up a Starter Kit's First SBR/Management Node | 99
- Setting Up the Second SBR/Management Node in a Starter Kit | 114
- Setting Up Data Node Hosts Included with the Starter Kit | 126
- Initial Cluster Startup Procedure | 132
- Recommendations before Configuring the Cluster | 137

This chapter documents the **configure** script that installs Session State Register software on the four node hosts in a Starter Kit cluster. If you have additional nodes to install beyond those four, you can specify them during the tasks presented in this chapter, although the text assumes the basic four nodes.

These topics are in the chapter:

## Cluster Installation Overview

Setting up the four nodes in a Session State Register Starter Kit is straightforward if you have performed the preparatory and planning steps and set up each server in order.

**NOTE:** Starting from SBR Release 7.6.0, an upgrade license is available to enable users who have purchased the Restricted Cluster Session license to upgrade to a Regular Cluster license. The Restricted Cluster license imposes a limit on the maximum number of concurrent sessions and does not allow addition of expansion kit license.

The workflow from this point includes these steps:



1. Unpack the SSR cluster software package on each node.
2. On the server you want to host the primary SBR Carrier and management nodes:
  - a. Run the configure script to specify the cluster topology and generate a set of baseline configuration files.
  - b. Run the configure script a second time to complete the software installation using the newly generated set of baseline configuration files.
3. On each of the other servers you want to host nodes:
  - a. Copy the baseline configuration files from the first server.
  - b. Run the installation script to load the software as specified in the configuration files
4. Start the cluster nodes in the correct order and monitor startup.
5. Set up the basic cluster configuration.

## Unpacking Session State Register Software

This section assumes that you are installing Steel-Belted Radius Carrier on your Solaris server for the first time or that you are installing Steel-Belted Radius Carrier in a directory other than the one used by previous installations—a clean installation.

These instructions may be used for unpacking software on any type of node host. Perform this step on each node host in the cluster before proceeding to the next step.

### Package Management Commands

The installer for the Solaris version of the Steel-Belted Radius Carrier server software uses **pkg** files with filenames that include the edition and version of the server software.

[Table 18 on page 95](#) lists common Solaris package management commands.

**Table 18: Useful Package Management Commands**

Command	Function
<code>pkginfo -x  grep "JNPR RSAR"</code>	Report any pre-existing packages and patches.
<code>pkginfo -l JNPRsbr</code>	Report high level description for specified package
<code>pkginfo -r JNPRsbr</code>	Show installed directory

Table 18: Useful Package Management Commands (*continued*)

Command	Function
<b>pkgadd -d / path/to /JNPRsbr.pkg [-a none] JNPRsbr.pkg</b>	Installs [with user prompts].  <b>NOTE:</b> You can install SBR at any location. We recommend you to install SBR at the default location "/opt/JNPRsbr".
<b>pkgrm JNPRsbr</b>	Uninstall Steel-Belted Radius Carrier.

To install the Steel-Belted Radius Carrier server software:

On the Solaris server:

1. Log in to the Solaris server as root.
2. Copy the Steel-Belted Radius Carrier installation files from their download location to the Solaris server. Make sure to copy them to a local or remote hard disk partition that is readable by root.

This example copies the files from a download directory to the **/tmp/sbr** directory.

Execute:

```
mkdir -p /opt/tmp  
cp -pR /tmp/sbr/solaris/* /opt/tmp/
```

3. Extract the Steel-Belted Radius Carrier installation package.

For 64-bit Solaris, execute:

```
cd /tmp/sbr  
ls -ltr
```

```
total 216240  
-rw-r--r--  1 root    root      110712276 Aug 25 09:44  
sbr-cl-8.6.0.R-1.sparcv9.tgz
```

Execute:

```
gunzip sbr-cl-8.6.0.R-1.sparcv9.tgz  
tar xf sbr-cl-8.6.0.R-1.sparcv9.tar
```

4. Verify that the extraction worked and confirm the name of the package file.

For 64-bit Solaris, execute:

```
ls -ltr
```

```
total 216256
drwxr-xr-x   4 Xtreece  other      370 Aug 24 17:01 JNPRsbr.pkg
-rw-r--r--   1 root     root       110712276 Aug 25 09:44
sbr-cl-8.6.0.R-1.sparcv9.tar
```

5. Install the package.

Execute:

**pkgadd -d /tmp/sbr**

```
The following packages are available:
1  JNPRsbr.pkg          JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier
Cluster Edition)          (sparc) 8.60.50006
Select package(s) you wish to process (or 'all' to process all packages).
(default: all) [?,??,q]:  all
```

6. Type **all** and press Enter.

The script resumes.

```
Processing package instance <JNPRsbr.pkg> from </tmp>
```

7. Confirm the installation directory.

Depending on the system configuration, the script prompts you to create the **/opt/JNPRsbr** directory if it does not exist, over-write an already-extracted package, or any of several other questions.

```
The selected base directory </opt/JNPRsbr> must exist before installation is
attempted.
```

```
Do you want this directory created now [y,n,?,q]
```

Answer the question appropriately (or change the extraction path if necessary) so that the script can proceed.

To accept the default directory as a target, enter **y**.

The script resumes.

```
Using </opt/JNPRsbr> as the package base directory.
#Processing package information.
#Processing system information.
    48 package pathnames are already properly installed.
```

```
#Verifying disk space requirements.
#Checking for conflicts with packages already installed.
#Checking for setuid/setgid programs.
```

```
This package contains scripts which will be executed with super-user
permission during the process of installing this package.
```

```
Do you want to continue with the installation of <JNPRsbr> [y,n,?]
```

8. Confirm that you want to continue to install the package. Enter **y**.

```
Installing JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster
Edition) as <JNPRsbr>
```

```
## Executing preinstall script.
## Installing part 1 of 1.
.
.
.
[ verifying class <none> ]
## Executing postinstall script.
Newly installed server directory will be backed up as:
/opt/JNPRsbr/radius/install/backups/2009:03:31-00:34:06
```

```
Installation of <JNPRsbr> was successful.
```

On the Linux server:

1. Log in to the Linux server as root.
2. Copy the Steel-Belted Radius Carrier installation files from their download location to the Linux server. Make sure to copy them to a local or remote hard disk partition that is readable by root.

This example copies the files from a download directory to the **/tmp/sbr** directory.

Execute:

```
mkdir -p /opt/tmp  
cp -pR /tmp/sbr/linux/* /opt/tmp/
```

3. Execute `cd /opt/tmp/`.

4. Install the package.

To install the 64-bit SBR package, execute:

`sbr-cl-8.6.0.R-1.el7.x86_64.rpm` on RHEL 7.3 or later.

See [“Linux” on page 47](#) for supported RHEL versions.

5. Installation proceeds and SBR is installed at `/opt/JNPRsbr/radius/install`.

## Setting Up a Starter Kit's First SBR/Management Node

This section describes installing software on the first server in a Starter Kit cluster, usually a server that you want to host a SBR node/management node (sm). The procedure for this first server is unique because it includes creating configuration files for all nodes in the cluster.

**TIP:** If you are going to use centralized Configuration Management (CCM) to share SBR Carrier configuration files between SBR nodes, remember that the primary server must be installed before replicas. You want the first server you work on in the cluster to be the CCM primary server.

### Configuring the Host Software on the First Server in the Cluster

Before starting this procedure, review [“Before You Install Software” on page 39](#). In particular, review requirements for: [“Setting Up External Database Connectivity \(Optional\)” on page 53](#) and [“Installing the SIGTRAN Interface \(Optional\)” on page 55](#), as steps in this procedure require the server to be preconfigured for these capabilities.

To install software on the first server in a cluster, which you want to host a SBR Carrier and a Management Node:

1. As root, navigate to the directory where you installed the Steel-Belted Radius Carrier package. For information about directory in which Steel-Belted Radius Carrier package is installed, see [“Unpacking Session State Register Software” on page 95.](#)

Then, navigate to the **radius/install** subdirectory and run:

Execute:

```
cd /opt/JNPRsbr/radius/install/
```

2. Run the **configure** script:

Execute:

```
./configure
```

3. Review and accept the Steel-Belted Radius Carrier license agreement.

Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

4. From the menu of configuration tasks, enter **2** to specify “Generate Cluster Definition.”

```
Configuring SBR Software
```

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node MyNode_1
is not configured and processes are down, needs to be configured
-----
```

```
1. Unconfigure Cluster Node
   Not used when merely updating existing cluster definitions.
```

```
2. Generate Cluster Definition
   Creates new or updates existing cluster definitions.
   Modifies the shared directory but does not modify this node.
```

```
3. Configure Cluster Node
   To be preceded by 'Generate Cluster Definition' on any node.
   Must be invoked on each and every node of the cluster.
```

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

Enter number of desired configuration task, or q to quit [2]: 2

#### 5. Specify the name of the cluster.

Enter the name exactly as you specified it in [Table 9 on page 35](#).

**Enter SBR cluster name [MyCluster]: MyCluster**

#### 6. Enter SSR license numbers and the number and type of nodes in the cluster.

If a Restricted Cluster license is provided as input for a Starter Kit license, you are not prompted for an Expansion Kit license. By default, you have a minimal cluster of two SBR nodes, two management nodes, and two data nodes.

The SBR Cluster Starter Kit license allows you to create a minimal cluster of 2 SBR nodes, 2 management nodes, and 2 data nodes. When each node is installed on a separate machine the cluster topology is denoted as {2s,2m,2d}. When SBR nodes are paired with management nodes on the same machines the cluster topology is denoted as {2sm,2d}.

An optional SBR Cluster Management Expansion Kit allows you to add a third management node for {2sm,1m,2d} and an optional Data Expansion Kit allows you to add 2 more data nodes for {2sm,1m,4d} clusters. Additional SBR licenses allow you to add up to 18 more SBR nodes to obtain a maximal cluster {18s,2sm,1m,4d} and/or enable extra features.

While it is not difficult to add management and/or SBR nodes to an existing cluster, adding data nodes is more difficult and may require you to shutdown the entire cluster as opposed to a rolling restart.

```
Enter Starter Kit license: 1234 5678 9100 1234 5678 9100 0050
Enter Management Expansion Kit license, if any:
Enter Data Expansion Kit license, if any:
Enter total number of SBR nodes to be configured [2]: 2
Enter number of management nodes to be paired with SBR nodes [2]: 2
```

7. Verify the configuration that you specified is accurate, and enter **y** to continue.

```
Creating cluster cambridge{0s,2sm,0m,2d}
will require 4 machines total. Do you wish to continue? [y]:y
```

```
All cluster nodes will share the same Session State Register (SSR).
Setting password for SSR admin account hadmsql
Password:
Again:
Setting password for SSR software account hadmsbr
Password:
Again:
```



Information will now be gathered for each machine in the cluster.  
You will have a chance to review all information at least once  
before any machines are modified.

8. Enter, for each node:

- The node name
- Type of node
- Node ID
- SBR Carrier license numbers (if required)
- The IP address for each node

The information you need is in [Table 9 on page 35](#).

Information will now be gathered for each machine in the cluster.  
You will have a chance to review all information at least once  
before any machines are modified.

```
-----
SBR 8.60.50006 cluster MyCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node MyNode_1
Partial configuration at present is {0s,0sm,0m,0d} of {0s,2sm,0m,2d}
-----
```

```
Enter node name [MyNode_1]: MyNode_1
Enter node type (sm) [sm]: sm
Enter SBR node ID (30-59) [30]:
Enter SBR node IP address by which it is known to management nodes.
Enter SBR node IP address: 192.168.0.1
Enter SBR licenses meant only for this particular SBR node.
Enter one license per line and an empty line when finished.
Enter SBR full license: xxxx xxxx xxxx xxxx
Enter SBR feature license:
Enter MGMT node ID (1-3) [1]:
Enter MGMT node IP address by which it is known to other nodes.
Enter MGMT node IP address: 192.168.0.1
```

```
-----
SBR 8.60.50006 cluster MyCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node MyNode_1(sm)
Partial configuration at present is {0s,1sm,0m,0d} of {0s,2sm,0m,2d}
-----
```

```
Enter node name [MyNode_2]: MyNode_2
Enter node type (sm,d) [sm]: sm
Enter SBR node ID (30-59) [31]:
Enter SBR node IP address by which it is known to management nodes.
Enter SBR node IP address: 192.168.0.2
Enter SBR licenses meant only for this particular SBR node.
Enter one license per line and an empty line when finished.
Enter SBR full license: xxxx xxxx xxxx xxxx
Enter SBR feature license:
Enter MGMT node ID (1-3) [2]:
Enter MGMT node IP address by which it is known to other nodes.
Enter MGMT node IP address: 192.168.0.2
-----
```

```
SBR 8.60.50006 cluster MyCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node MyNode_2(sm)
Partial configuration at present is {0s,2sm,0m,0d} of {0s,2sm,0m,2d}
-----
```

```
Enter node name [MyNode_3]: MyNode_3
Enter node type (d) [d]:
Enter DATA node ID (10-19) [10]: 18
Enter DATA node IP address by which it is known to management nodes.
Enter DATA node IP address: 192.168.0.18
-----
```

```
SBR 8.60.50006 cluster MyCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node MyNode_3 (sm)
Partial configuration at present is {0s,2sm,0m,1d} of {0s,2sm,0m,2d}
-----
```

```
Enter node name [MyNode_4]: MyNode_4
Enter node type (d) [d]:
Enter DATA node ID (10-19) [11]: 19
Enter DATA node IP address by which it is known to management nodes.
Enter DATA node IP address: 192.168.0.19
-----
```

```
SBR 8.60.50006 cluster MyCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node MyNode_4
Generated configuration is {0s,2sm,0m,2d} of {0s,2sm,0m,2d}
-----
```

9. The system generates the required configuration files and prompts you to view, accept, or reject them. Enter **a** to accept them and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

```
Reviewing Configuration Files...
```

```
-----
```

```
/opt/JNPRsbr/radius/install/tmp/dbcluster.rc
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
```

```
View (v), accept (a), or reject (r) configuration files: a
```

10. From the menu of configuration tasks, enter **3** to specify “Configure Cluster Node.”

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node MyNode_2
is not configured and processes are down, needs to be configured
-----
```

```
1. Unconfigure Cluster Node
   Not used when merely updating existing cluster definitions.
```

## 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.

## 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on any node.  
Must be invoked on each and every node of the cluster.

## 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

## 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

## 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter number of desired configuration task, or q to quit [2]: **3**

## 11. Specify the name of the cluster.

Enter the name exactly as you specified it in [Table 9 on page 35](#).

Enter SBR cluster name [MyCluster]: **MyCluster**

12. Specify whether you want to use the JRE installed in your system to enable JDBC plug-ins and JavaScript implementation.

```
Do you want to configure Java Runtime Environment for JDBC Feature [n] :
```

- If no, press Enter to proceed to the next prompt. SBR Carrier does not support JDBC plug-ins unless you specify a valid JRE path.
- If yes, type **y** and press Enter. You are prompted to specify the path where the JRE is installed in your system. The Java Virtual Machine (JVM) architecture should be compatible with SBR Carrier.

**NOTE:** Java 1.8.0 or a later version is required to access the Web GUI. To support both JDBC plug-ins and Web GUI, it is recommended to use Java 1.8.0 or a later version with the JVM architecture compatible with your SBR Carrier. For example, if you are using the 64-bit version of SBR Carrier, you must use the 64-bit version of Java 1.8.0 or later.

```
Enter 64-bit libjvm.so path (Ex: /opt/jvm/jre/lib/amd64/server/ ) :
```

**NOTE:** If you enter an incorrect JVM path three times, SBR Carrier proceeds to the next step. In this case, you will not be able to use JDBC plug-ins. To specify the valid JVM path, you need to run the configure script again.

13. The script prompts for the type of installation, either a new installation or a migration from an earlier release:

```
Please enter backup or radius directory from which to migrate.
Enter n for new configuration, s to search, or q to quit
[n]: n
```

- For a new installation, enter **n**.
- If you are migrating an existing Steel-Belted Radius Carrier installation and have copied a previous release's files to the Release 8.6.0 server (in [“Creating a Copy of Existing SBR Server Release Files for Migration” on page 54](#)), enter the directory path to the copy of the old installation.

For example:

```
[ /opt/JNPRsbr/PreviousInstall]:
```

- If you want to search well known locations in the file system for pre-existing installations or backups of SBR software, enter **s**. You are presented with a list of directories from which you may choose to migrate, if any are found.

14. Specify the username of the UNIX account that you intend to use to gain access to this machine using the Web GUI. The UNIX account must be defined on this machine (for example by **/etc/passwd** or **NIS**) and must have a valid, non-empty password. Additional accounts may be specified using the Web GUI.

```
Enter initial admin user (UNIX account must have a valid password) [root]:
```

Press Enter to accept the default, root.

15. Specify whether you want to set up Centralized Configuration Management (CCM).

CCM allows you to replicate substantial portions of a particular SBR node's configuration across multiple SBR nodes. The SBR node to be replicated is known as the *primary* and the SBR nodes on which the primary configuration is replicated are known as *replicas*. The entire set of an SBR primary and all of its replicas are known as an *SBR replication group*. Although any given SBR replication group typically includes all SBR nodes in a particular SSR cluster, an SBR replication group can span multiple clusters (including standalone SBR nodes) or can span only a few SBR nodes in a cluster. SBR nodes for which CCM is not enabled are known as *autonomous* because they are configured independently of one another.

The rest of this task assumes that you want to configure a typical SBR replication group that includes all SBR nodes in this cluster. If you do not want to enable CCM then skip the rest of this task.

```
Enable Centralized Configuration Management (CCM) for this SBR node? [n]: y
```

16. Specify the server role.

- a. When CCM is enabled, configure this first SBR/management (sm) node as an SBR primary.

```
Configure SBR node as CCM primary (p) or replica (r)? [r]: p
```

- b. You are prompted to configure the primary secret you want to use on this SBR replication group:

```
Enter primary host secret:
```

Type the secret string and press Enter.

- c. You are prompted to confirm the primary secret.

**Confirm primary host secret:**

Type the secret string again and press Enter.

17. Specify whether you want to use the auto-restart module that automatically restarts the SBR Carrier server in case of an unexpected shutdown.

```
Do you want to enable "Radius WatchDog" Process? [n]: Y
Radius WatchDog feature set to Enable
Please ensure that Perl 5 or better is installed.
```

**NOTE:** If Perl version 5 is not installed, the radiusd script will not run, even if enabled by configuration, and SBR Carrier will operate without the auto-restart module running.

18. Specify whether you want to configure SBR Carrier to provide LDAP server emulation for configuration and statistics using the LCI.

```
Do you want to enable LCI? [n] :
```

- If no, press Enter to accept the default.
- If yes, enter **y** and press Enter. You are prompted to provide information for LCI configuration.
  - a. When you are prompted for the port number, enter the port number that is used for communication between SBR Carrier and the LDAP client.

**NOTE:** SBR Carrier uses port 667 as the default for LDAP emulation to avoid conflict with other LDAP servers.

```
Configure LCI Port [667]: 1026
```

- b. The script displays the interfaces available in the system. When you are prompted to enter interface addresses on which Steel-Belted Radius Carrier should listen for LCI requests, enter the addresses you want to use from the Available Interfaces list.

```
LCI Interface Configuration :
Available interfaces :
127.0.0.1
10.212.10.66

HELP : Enter one interface per line and an empty line when finished.

Enter LCI interface addresses from the above list.

Enter LCI interface address : 10.212.10.66
Enter LCI interface address : 127.0.0.1
Enter LCI interface address :
```

**NOTE:** SBR Carrier uses all interfaces for listening to LCI requests if you do not enter any interfaces.

- c. Specify whether you want to change the default LCI password to prevent unauthorized LDAP clients from accessing your database.

```
Do you want to change LCI Password? [n]:
```

- If no, press Enter to accept the default password.
- If yes, enter y and press Enter. You are prompted to enter a new password.

```
Do you want to change LCI Password? [n]: Y

Password must meet the following requirements:

1. 6-8 Alphanumeric characters.
2. No Special characters other than underscore ('_').

Enter Password:
Confirm Password:
Password will be changed when SBR restarts.
```



**NOTE:** Make sure that the entered password is at least 6 alphanumeric characters and not more than 8 characters in length. The password should not include any special characters other than underscore ('\_').

**NOTE:** The configure script also checks whether the LDAP utilities (such as **ldapdelete**, **ldapmodify**, and **ldapsearch**) are installed in your system. For Linux, a warning message is displayed if you have not installed any of these utilities in your system. For Solaris, LDAP utilities are shipped with SBR Carrier package.

19. Specify whether you want to configure Steel-Belted Radius Carrier for use with an Oracle database.

To support this option, the server must already be configured as an Oracle client. (See [“Setting Up External Database Connectivity \(Optional\)” on page 53.](#))

**Configuring for use with generic database**

**Do you want to configure for use with Oracle? [n]:**

If no, press Enter to accept the default.

If yes, type **y** and press Enter. You are prompted for version and path information for the Oracle library files.

**Do you want to configure for use with Oracle? [n]: y**

**Supported Oracle versions: 10, 11, 12**

**What version of Oracle will be used? [10]: 10**

**Configuring for use with Oracle 10**

**Setting the environment variable ORACLE\_HOME**

**Enter ORACLE\_HOME [/dbms/u10/app/oracle/product/10.2.0]:**

**Setting the environment variable LD\_LIBRARY\_PATH**

**Enter path for Oracle shared libraries**

**[/dbms/u10/app/oracle/product/10.2.0/lib]**

**Setting the environment variable TNS\_ADMIN**

**Enter TNS\_AMDIN [/dbms/u10/app/oracle/product/10.2.0/network/admin]:**

**NOTE:** You must configure 64-bit Oracle client for 64-bit SBR Carrier.

20. Specify whether you want the Steel-Belted Radius Carrier server to communicate with an SS7 system using SIGTRAN.

If your SBR Carrier is using the optional SIM authentication module, or the WiMAX module with the EAP-AKA protocol, or the server is interfacing with a UMA or Femtocell network, you need to have Signalware installed to communicate with the SS7 network.

To support this option, the server must already be configured to support SIGTRAN using Signalware. (See “[Installing the SIGTRAN Interface \(Optional\)](#)” on page 55 for an overview, and “[SIGTRAN Support for Steel-Belted Radius Carrier](#)” on page 309 for specific instructions.)

```
Do you want to configure for use with SIGTRAN? [n]: y
Configuring for use with SIGTRAN
Setting the environment variable OMNI_HOME
Enter OMNI_HOME [/opt/JNPRss7]:
```

21. Specify whether you want to start the GWrelay process while executing the `./sbrd start` script.

```
Do you want to enable "GWrelay" Process? [n]: y
GWrelay will be started with sbrd
```

22. Specify whether you want to install the optional SNMP module to monitor your Steel-Belted Radius Carrier server from an SNMP management station.

**Do you want to configure SNMP? [n]:**

- If no, press Enter to proceed to the next prompt.
- If yes, type **y** and press Enter. The installer prompts you for the information it needs to configure the `jnpnsnmpd.conf` and `startsnmp.sh` files.
  - a. When you are prompted for a community string, enter the community string used to validate information sent from the SNMP subagent on the Steel-Belted Radius Carrier server to your SNMP management station.

**Choose a community string: public**

- b. When you are prompted for a range of IPv4 addresses, specify a starting IP address in Classless Inter-Domain Routing format. To specify that only one host may query the agent, enter the IP address of the host followed by `/32`. To specify that any host on a designated class C network may query the agent, enter the starting address of the network followed by `/24`.

**Specify the range of IPv4 addresses that may query this agent, such as 1.2.3.0/24.**

**Address range: 192.168.70.0/24**

- c. If you are using SNMPv2, enter the DNS name or IP address of the SNMP management station to receive trap information from the SNMP subagent on the Steel-Belted Radius Carrier server.

**SNMPv2 trap sink: 192.168.70.86**

- d. Set the SNMP agent port.

Although you may specify the default SNMP port, 161, we recommend that you specify a different port to avoid contention with other agents that are likely to already be using 161. If you choose an alternate port, make a note of it because your SNMP management station needs to be configured to the same setting.

**Specify SNMP agent listening port[161]: 24161**

- e. Specify a trap sink address, if required.

**Optionally specify a trap sink address that will receive SNMPv2 trap**

**[localhost]: 172.28.72.83 2**

**SNMPv2 trap sink port[162]:**

**Configuration of SNMP complete.**

23. The script searches for the Java 1.8.0 or later version in the default system path and displays a confirmation message if found.

```
Configuring Admin GUI Webserver
Compatible Java version 1.8.0_66 found in: /usr/java/jdk1.8.0_66
```

If the specific version is not found, the script prompts you to enter the path where the specific Java version is installed in your system.

```
Enter Java version 1.8 installed path :
```

24. Specify whether you want to install a custom SSL certificate for the Web GUI.

```
Do you want to install custom SSL certificate for Admin WebServer? [n]:
```

- If no, press Enter. A self-signed certificate is created and installed in your server.
- If yes, enter **y** and press Enter. You are prompted to enter the absolute path where the SSL certificate is available. For example, **/opt/customSSLCert.pfx**.

```
Enter the absolute path to certificate.
Note: Only *.pfx files are accepted. (Example-/opt/customSSLCert.pfx):
```

When you are prompted for the password, enter the password to open the SSL certificate.

Enter the password to open the certificate :

20. Specify whether you want to configure the Steel-Belted Radius Carrier software (both RADIUS and SSR processes as appropriate for the given node type) to autoboot (restart automatically when the operating system is restarted). We recommend that you enable autoboot behavior.

**Enable (e), disable (d), or preserve (p) autoboot scripts [e]: e**

A local **radiusdir/radius/sbrd** script is always created, and **/opt/JNPRhadm/sbrd** is always a symbolic link to this local copy.

- If you enter e (enable), the **configure** script copies the local **sbrd** script to **/etc/init.d**, where it is automatically invoked by the OS whenever the OS is stopped or started.
- If you enter d (disable), the **configure** script removes all copies of the **sbrd** script from **/etc/init.d**, thus disabling autoboot for all versions of the Steel-Belted Radius Carrier.
- If you enter p (preserve), the **configure** script does nothing, thereby leaving your previous autoboot scripts unchanged.

When you finish entering settings, the script configures Steel-Belted Radius Carrier with the specified settings and then displays:

```
The SBR Admin Web GUI can be launched using the following URL:
https://<servername>:2909

Configuration complete
```

21. Configure the second SBR/management node in the Starter Kit by following the procedure described in [“Setting Up the Second SBR/Management Node in a Starter Kit” on page 114](#).

## Setting Up the Second SBR/Management Node in a Starter Kit

This section describes installing the software on the second SBR/management node host in a Starter Kit cluster, following the initial combined SBR/management node installation. (See [“Setting Up a Starter Kit’s First SBR/Management Node” on page 99](#).)



**CAUTION:** If you have not installed the Cluster’s first server yet, do not proceed with the cluster installation.

## Populating the JNPRShare Directory

Before running the configure script, make a local copy of the configuration files that were created during installation on the first server. (In the following procedures, FTP is used to copy files. You may use whatever distribution mechanism you want to keep the share directories synchronized, so long as all **/JNPRshare** directories on all servers are readable and writable by the root and hadm accounts.)

To copy the cluster's base configuration files to this target machine:

1. Log in as **hadm**.

2. Change directories to the install directory on the local server.

Execute:

```
cd /opt/JNPRshare/install/ <cluster_name>
```

For example:

```
cd /opt/JNPRshare/install/MyCluster
```

3. Use FTP binary mode to connect to the first server that was set up and navigate to the **radius/install** subdirectory of the directory in which the **JNPRsbr** package was installed (**/opt/JNPRsbr/radius/install** by default) on the source server.

4. Execute a **get** command to transfer the **configure <cluster name> .tar** file to the local directory.

For example:

```
bin
```

```
get configure.MyCluster.tar
```

5. Extract the configuration files from the archive.

For example:

```
tar xvf configure.MyCluster.tar
```

The output display includes five files similar to this example:

```
$ tar xvf configure.MyCluster.tar
x dbcluster.rc, 1925 bytes, 4 tape blocks
x config.ini, 2435 bytes, 5 tape blocks
x my.cnf, 1017 bytes, 2 tape blocks
x dbclusterndb.gen, 33474 bytes, 66 tape blocks
x dbcluster.dat, 921 bytes, 2 tape blocks
```

## Configuring Host Software on the Second SBR Carrier and Management Node

Before starting this procedure, review [“Before You Install Software” on page 39](#). In particular, review requirements for: [“Setting Up External Database Connectivity \(Optional\)” on page 53](#) and [“Installing the](#)

[SIGTRAN Interface \(Optional\)](#) on page 55, as steps in this procedure require the server to be preconfigured for these capabilities.

To configure SSR software on the second SBR Carrier and management node host in a Starter Kit cluster:

1. As root, navigate to the directory where you installed the Steel-Belted Radius Carrier package. For information about directory in which Steel-Belted Radius Carrier package is installed, see [“Unpacking Session State Register Software”](#) on page 95.

Navigate to the **radius/install** subdirectory and run:

Execute:

```
cd /opt/JNPRsbr/radius/install/
```

2. Execute the **configure** script to install the Steel-Belted Radius Carrier server software:

Execute:

```
./configure
```

3. Review and accept the Steel-Belted Radius Carrier license agreement.

Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

4. From the menu of configuration tasks, enter **3** to specify **Configure Cluster Node**.

```
Configuring SBR Software
```

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node MyNode_2
is not configured and processes are down, needs to be configured
-----
```

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.

## 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.

## 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on any node.  
Must be invoked on each and every node of the cluster.

## 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

## 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

## 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter number of desired configuration task, or q to quit [2]: **3**

## 5. Specify the name of the cluster.

Enter the name exactly as you specified it in [Table 9 on page 35](#).

Enter SBR cluster name [MyCluster]: **MyCluster**

You are prompted to verify whether you want to proceed, unless the script detects any unusual installation conditions (a pre-existing directory, for example). In some cases, you may be prompted to resolve or ignore them.

6. The system reads the configuration files that you copied to the server and prompts you to change some settings from the original server's to adapt them to this server. Enter **y** to proceed.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? [n]: y
```

```
Cleaning old directories
/opt/JNPRhadm
-----
SBR 8.60.50006 cluster MyCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node MyNode_2(sm)
Reviewing Configuration Files...
-----
```

```
/opt/JNPRsbr/radius/install/tmp/dbcluster.rc
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
```

```
View (v), accept (a), or reject (r) configuration files: a
```

7. Enter **a** to accept the modified configuration files and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

8. Specify whether you want to use the JRE installed in your system to enable JDBC plug-ins and JavaScript implementation.



```
Do you want to configure Java Runtime Environment for JDBC Feature [n] :
```

- If no, press Enter to proceed to the next prompt. SBR Carrier does not support JDBC plug-ins unless you specify a valid JRE path.
- If yes, type **y** and press Enter. You are prompted to specify the path where the JRE is installed in your system. The Java Virtual Machine (JVM) architecture should be compatible with SBR Carrier.

**NOTE:** Java 1.8.0 or a later version is required to access the Web GUI. To support both JDBC plug-ins and Web GUI, it is recommended to use Java 1.8.0 or a later version with the JVM architecture compatible with your SBR Carrier. For example, if you are using the 64-bit version of SBR Carrier, you must use the 64-bit version of Java 1.8.0 or later.

```
Enter 64-bit libjvm.so path (Ex: /opt/jvm/jre/lib/amd64/server/ ) :
```

**NOTE:** If you enter an incorrect JVM path three times, SBR Carrier proceeds to the next step. In this case, you will not be able to use JDBC plug-ins. To specify the valid JVM path, you need to run the configure script again.

#### 9. Specify the server's initial admin user.

```
Enter initial admin user (UNIX account must have a valid password) [root]:
```

Press Enter to accept the default, root.

#### 10. Specify whether you want to set up Centralized Configuration Management (CCM).

```
Enable Centralized Configuration Management (CCM) for this SBR node? [n]: y
```

Then specify the server role.

```
Configure SBR node as CCM primary (p) or replica (r)? [r]: r
```

- a. Because this is not the first server to be installed, if CCM is used, this server cannot be the primary. Enter **r** to set the role as replica.
- b. Specify whether you want a local replica.ccmpkg file from the primary server.

```
Have you got a local replica.ccmpkg file from the primary (y/n)? [y]:
```

- c. Specify the primary CCM server's hostname.

**Enter primary host name:**

For example:

**Enter primary host name: MyNode\_1**

- d. Specify the primary CCM server's IP address.

**Enter primary host IPv4 address [172.28.84.35]:**

For example:

**Enter primary host IPv4 address [172.28.84.35]: 198.168.0.1**

- e. Specify the primary host secret string.

**Enter primary host secret:**

Type the secret string and press Enter.

- f. Confirm the host secret string.

**Confirm primary host secret:**

Type the secret string again and press Enter.

11. Specify whether you want to use the auto-restart module that automatically restarts the SBR Carrier server in case of an unexpected shutdown.

```
Do you want to enable "Radius WatchDog" Process? [n]: Y
Radius WatchDog feature set to Enable
Please ensure that Perl 5 or better is installed.
```

**NOTE:** If Perl version 5 is not installed, the radiusd script will not run, even if enabled by configuration, and SBR Carrier will operate without the auto-restart module running.

12. Specify whether you want to configure SBR Carrier to provide LDAP server emulation for configuration and statistics using the LCI.

```
Do you want to enable LCI? [n] :
```

- If no, press Enter to accept the default.
- If yes, enter y and press Enter. You are prompted to provide information for LCI configuration.
  - a. When you are prompted for the port number, enter the port number that is used for communication between SBR Carrier and the LDAP client.

**NOTE:** SBR Carrier uses port 667 as the default for LDAP emulation to avoid conflict with other LDAP servers.

```
Configure LCI Port [667]: 1026
```

- b. The script displays the interfaces available in the system. When you are prompted to enter interface addresses on which Steel-Belted Radius Carrier should listen for LCI requests, enter the addresses you want to use from the Available Interfaces list.

```
LCI Interface Configuration :
Available interfaces :
127.0.0.1
10.212.10.66

HELP : Enter one interface per line and an empty line when finished.

Enter LCI interface addresses from the above list.

Enter LCI interface address : 10.212.10.66
Enter LCI interface address : 127.0.0.1
Enter LCI interface address :
```

**NOTE:** SBR Carrier uses all interfaces for listening to LCI requests if you do not enter any interfaces.

- c. Specify whether you want to change the default LCI password to prevent unauthorized LDAP clients from accessing your database.

```
Do you want to change LCI Password? [n]:
```

- If no, press Enter to accept the default password.
- If yes, enter y and press Enter. You are prompted to enter a new password.

```
Do you want to change LCI Password? [n]: Y
```

```
Password must meet the following requirements:
```

1. 6-8 Alphanumeric characters.
2. No Special characters other than underscore ('\_').

```
Enter Password:
```

```
Confirm Password:
```

```
Password will be changed when SBR restarts.
```

**NOTE:** Make sure that the entered password is at least 6 alphanumeric characters and not more than 8 characters in length. The password should not include any special characters other than underscore ('\_').

**NOTE:** The configure script also checks whether the LDAP utilities (such as **ldapdelete**, **ldapmodify**, and **ldapsearch**) are installed in your system. For Linux, a warning message is displayed if you have not installed any of these utilities in your system. For Solaris, LDAP utilities are shipped with SBR Carrier package.

13. Specify whether you want to configure Steel-Belted Radius Carrier for use with an Oracle database.

To support this option, the server must already be configured as an Oracle client (See [“Setting Up External Database Connectivity \(Optional\)” on page 53.](#)).

**Configuring for use with generic database**

**Do you want to configure for use with Oracle? [n]:**

If no, press Enter.

If yes, type y and press Enter. You are prompted for version and path information for the Oracle library files.

**Do you want to configure for use with Oracle? [n]: y**

**Supported Oracle versions: 10, 11, 12**

What version of Oracle will be used? [10]: 10  
 Configuring for use with Oracle 10  
 Setting the environment variable ORACLE\_HOME  
 Enter ORACLE\_HOME [/dbms/u10/app/oracle/product/10.2.0]:  
 Setting the environment variable LD\_LIBRARY\_PATH  
 Enter path for Oracle shared libraries  
 [/dbms/u10/app/oracle/product/10.2.0/lib]:  
 Setting the environment variable TNS\_ADMIN  
 Enter TNS\_ADMIN  
 [/dbms/u10/app/oracle/product/10.2.0/network/admin]:

**NOTE:** You must configure 64-bit Oracle client for 64-bit SBR Carrier.

14. Specify whether you want the Steel-Belted Radius Carrier server to communicate with an SS7 system using SIGTRAN.

To support this option, the server must already be configured to support SIGTRAN using Signalware. (See [“Installing the SIGTRAN Interface \(Optional\)” on page 55](#) for an overview, and [“SIGTRAN Support for Steel-Belted Radius Carrier” on page 309](#) for specific instructions.)

```
Do you want to configure for use with SIGTRAN? [n]: y
Configuring for use with SIGTRAN
Setting the environment variable OMNI_HOME
Enter OMNI_HOME [/opt/JNPRss7]:
```

15. Specify whether you want to start the GWrelay process while executing the `./sbrd start` script.

```
Do you want to enable "GWrelay" Process? [n]: y
GWrelay feature set to Enable
```

16. Specify whether you want to install the optional SNMP module to monitor your Steel-Belted Radius Carrier server from an SNMP management station.

**Do you want to configure SNMP? [n]:**

- If no, press Enter to proceed to the next prompt.
- If yes, type **y** and press Enter. The installer prompts you for the information it needs to configure the `jnprsnmpd.conf` and `startsnmp.sh` files.

- a. When you are prompted for a community string, enter the community string used to validate information sent from the SNMP subagent on the Steel-Belted Radius Carrier server to your SNMP management station.

**Choose a community string: public**

- b. When you are prompted for a range of IPv4 addresses, specify a starting IP address in Classless Inter-Domain Routing format. To specify that only one host may query the agent, enter the IP address of the host followed by /32. To specify that any host on a designated class C network may query the agent, enter the starting address of the network followed by /24.

**Specify the range of IPv4 addresses that may query this agent, such as 1.2.3.0/24.**

**Address range: 192.168.70.0/24**

- c. If you are using SNMPv2, enter the DNS name or IP address of the trap sink that you want to receive trap information from the SNMP subagent on the Steel-Belted Radius Carrier server.

**SNMPv2 trap sink: 192.168.70.86**

- d. Set the SNMP agent port.

Although you may specify the default SNMP port, 161, we recommend that you specify a different port to avoid contention with other agents that are likely to already be using 161. If you choose an alternate port, make a note of it because your MIB browser needs to be configured to the same setting.

**Specify SNMP agent listening port[161]: 24161**

- e. Specify a trap sink address, if required.

**Optionally specify a trap sink address that will receive SNMPv2 trap**

**[localhost]: 172.28.72.83 2**

**SNMPv2 trap sink port[162]:**

**Configuration of SNMP complete.**

17. The script searches for the Java 1.8.0 or later version in the default system path and displays a confirmation message if found.

```
Configuring Admin GUI Webserver
Compatible Java version 1.8.0_66 found in: /usr/java/jdk1.8.0_66
```

If the specific version is not found, the script prompts you to enter the path where the specific Java version is installed in your system.

```
Enter Java version 1.8 installed path :
```

18. Specify whether you want to install a custom SSL certificate for the Web GUI.

```
Do you want to install custom SSL certificate for Admin WebServer? [n]:
```

- If no, press Enter. A self-signed certificate is created and installed in your server.
- If yes, enter **y** and press Enter. You are prompted to enter the absolute path where the SSL certificate is available. For example, **/opt/customSSLCert.pfx**.

```
Enter the absolute path to certificate.  
Note: Only *.pfx files are accepted. (Example-/opt/customSSLCert.pfx):
```

When you are prompted for the password, enter the password to open the SSL certificate.

```
Enter the password to open the certificate :
```

14. Specify whether you want to configure the Steel-Belted Radius Carrier software (both RADIUS and SSR processes as appropriate for the given node type) to autoboot (restart automatically when the operating system is restarted). We recommend that you enable this behavior.

**Enable (e), disable (d), or preserve (p) autoboot scripts [e]: e**

A local **/radiusdir/radius/sbrd** script is always created, and **/opt/JNPRhadm/sbrd** is always a symbolic link to this local copy.

- If you enter **e** (enable), the **configure** script copies the local **sbrd** script to **/etc/init.d**, where it is automatically invoked by the OS whenever the OS is stopped or started.
- If you enter **d** (disable), the **configure** script removes all copies of the **sbrd** script from **/etc/init.d**, thus, disabling autoboot for all versions of Steel-Belted Radius Carrier.
- If you enter **p** (preserve), the **configure** script does nothing, thereby leaving your previous autoboot scripts unchanged.

15. When prompted, confirm that you want the installation to proceed and to configure Steel-Belted Radius Carrier with the specified settings. When the script finishes, it displays:

```
The SBR Admin Web GUI can be launched using the following URL:  
https://<servername>:2909  
  
Configuration complete
```

16. Configure the two data nodes in the Starter Kit using the procedure described in [“Setting Up Data Node Hosts Included with the Starter Kit”](#) on page 126.

## Setting Up Data Node Hosts Included with the Starter Kit

Use this procedure for any data node host installation. The examples in this section install and set up each of the two data nodes in a Starter Kit, using the **MyCluster** cluster example, but the procedure is the same for expansion kit installations.

**NOTE:** For performance reasons, the SSR ndbmt processes on data (D) nodes are configured to execute under the UNIX root account by default, as opposed to the UNIX hadm account. In particular, this allows the ndbmt processes to lock data in physical memory, which is faster, as opposed to allowing the OS to use swap space on disk, which is slower. The UNIX root account privilege is required in order to lock data in physical memory.

- The relevant configuration item is the **#sbrd-ndbd-run-as-root = true** parameter in the **[ndbd]** section of the **/opt/JNPRhadm/my.cnf** file. Note that the leading # character is required to distinguish this parameter as a **sbrd** script parameter; this parameter is *not* a comment and is always active. When the value of this parameter is true, the ndbmt processes execute under the UNIX root account. When the value of this parameter is false (or if the parameter is missing entirely), the ndbmt processes execute under the UNIX hadm account. The value of this parameter can only be changed immediately after configuring a data (D) node. The value of this parameter cannot be changed after the SSR processes are running.
- We recommend, although it is not necessary, that the parameter be configured the same on all data (D) nodes. In order to change the value of this parameter at a later time, you must unconfigure the data (D) node and then reconfigure it again.
- When the ndbmt processes are executed under the UNIX root account, it is extremely important that the **DataMemory** parameter in the **[ndbd default]** section of the **/opt/JNPRhadm/config.ini** file be configured properly with respect to the amount of physical memory that is actually available on the data (D) node. If the data (D) node does not have enough physical memory available, then the ndbmt processes can starve the entire machine, including the OS itself, for memory. By default, SBRC is configured under the assumption that at least 8 GB of memory is available *solely* for ndbmt processes. In practice, more than 8 GB is required to support the OS and other applications.

### Striping Data Nodes

For performance reasons, the data stored in the Session State Register (SSR) should be striped. If you choose not to enable striping, the SBR Carrier software operates in demonstration mode without enforcing minimum memory requirements. When operating in demonstration mode, the SBR Carrier software makes a best effort attempt to operate in spite of various deficiencies that would normally prevent operation due to poor performance.



The choice of whether or not to stripe must be answered when the `./configure` script (typically found in the `/opt/JNPR sbr/radius/install` directory) is executed in order to create a new cluster definition. When you execute the `./configure` script, and select option 2, **"Generate Cluster Definition"**, you are presented with the following prompts:

```
...
Enter number of management nodes to be paired with SBR nodes [2]: 2

Your license allows striping sun4v class hardware for performance.
However, striping requires at least 8GB memory on all data nodes.
The software will operate in demonstration mode with degraded
performance if you do not enable striping.  Enable striping? [y]: y

Creating cluster MyCluster{0s,2sm,0m,2d}
will require 4 machines total.  Do you wish to continue? [y]: y
...
```

If the prompts related to striping are not answered correctly (for example, striping is enabled but one or more data nodes has less than 8 GB memory, then you will not be able to configure all of the data nodes. In this case, when you execute the `./configure` script, and select option 3, **"Configure Cluster Node"**, and then select the **(c) Create** option, you are prompted as follows:

```
...
Create (c) new or update (u) existing node configuration? [u]: c
...
WARNING: d nodes require at least 8 GB physical memory
        whereas this machine has only 4 GB installed.
ERROR: Insufficient hardware
HINT: You may wish to reconfigure for demonstration mode instead.
...
```

The number of stripes is presently a fixed parameter, always being set to either 1 (striping disabled for demonstration mode), 4 (striping enabled for cluster), or 8 (striping enabled for standalone server or transition server). After the number of stripes is configured, it cannot be changed without destroying and then re-creating the entire SBR Carrier cluster. Again, because striping is a global parameter with respect to cluster geometry, all data nodes must always have the same number of stripes.

Each stripe is implemented by a separate SSR data process requiring its own unique node ID. Thus, eight node IDs are required for each data node in a transition server when striping is enabled. However, the `./configure` script only prompts for one base node ID per data node regardless of whether striping is enabled because higher order node IDs are determined by an algorithm related to the number of data

nodes and the number of stripes. Also, the `./sbrd` script (typically found in `/opt/JNPRsbr/radius`) operates upon all of the SSR data processes on a particular node as if they were one entity.

If any SSR data processes diverge from the group, the `./sbrd` script may detect this and warn you if you attempt to restart them. (You are not likely to encounter this unless you are having trouble starting the software in the first place.)

```
sbrd: WARNING: some ssr data processes failed, stop the survivors first
```

If you see this warning, use the `./sbrd status` command to verify whether any data processes have failed. If any data processes have failed while other data processes still persist, then execute `./sbrd stop ssr` followed by `./sbrd start ssr` and finally `./sbrd status` again to verify that the problem has been resolved.

When `./sbrd status` is executed as either root or hadm on a running M or SM node, or on a cluster that is striped, you should observe four times as many `[ndbd(NDB)]` nodes as there are actual data nodes (because four is the number of stripes). When `./sbrd status` is executed on a running data node, you should observe twice as many `ndbmtd` processes (the SSR data processes) as stripes because each working `ndbmtd` process is paired with a watchdog instance of itself to guard against failure.

## Populating the JNPRShare Directory

Before running the **configure** script, make a local copy of the configuration files that were created during installation on the first combined.

To copy the cluster's base configuration files to this target machine:

1. Log in as **hadm**.
2. Change directories to the working directory on the local server.

Execute:

```
cd /opt/JNPRshare/install/ <cluster_name>
```

For example:

```
cd /opt/JNPRshare/install/MyCluster
```

3. Use FTP binary mode to connect to the first server that was set up and navigate to the **radius/install** subdirectory of the directory in which the **JNPRsbr** package was installed (`/opt/JNPRsbr/radius/install` by default) on the source server.
4. Execute a **get** command to transfer the **configure. <cluster name> .tar** file to the local directory.

For example:

```
bin
```

```
get configure.MyCluster.tar
```

5. Extract the configuration files from the archive.

For example:

**tar xvf configure.MyCluster.tar**

The output display includes five files similar to this example:

```
$ tar xvf configure.cambridge.tar
x dbcluster.rc, 1925 bytes, 4 tape blocks
x config.ini, 2435 bytes, 5 tape blocks
x my.cnf, 1017 bytes, 2 tape blocks
x dbclusterndb.gen, 33474 bytes, 66 tape blocks
x dbcluster.dat, 921 bytes, 2 tape blocks
```

## Configuring the Host Software on the Data Nodes

To configure the software on a data node in a Starter Kit cluster:

**NOTE:** You must repeat this procedure on every data node in the cluster.

1. As root, navigate to the directory where you installed the Steel-Belted Radius Carrier package. For information about directory in which Steel-Belted Radius Carrier package is installed, see [“Unpacking Session State Register Software” on page 95](#).

Navigate to the **radius/install** subdirectory and run:

Execute:

**cd /opt/JNPRsbr/radius/install/**

2. Execute the **configure** script to install the Steel-Belted Radius Carrier server software:

Execute:

**./configure**

3. Review and accept the Steel-Belted Radius Carrier license agreement.

Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

4. From the menu of configuration tasks, enter **3** to specify **Configure Cluster Node**.

Configuring SBR Software

```

-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node MyNode_3
is not configured and processes are down, needs to be configured
-----

```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on any node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter number of desired configuration task, or q to quit [2]: **3**

#### 5. Specify the name of the cluster.

Enter the name exactly as you specified it in [Table 9 on page 35](#).

Enter SBR cluster name [MyCluster]: **MyCluster**

You are prompted to verify whether you want to proceed, unless the script detects any unusual installation conditions (a pre-existing directory, for example). In some cases, you may be prompted to resolve or ignore them.

#### 6. The system reads the configuration files that you copied to the server and prompts you to change some settings to adapt them to this server. Enter **y** to proceed.

WARNING: You are about to make irreversible changes to this node.  
Are you sure that you wish to continue? [n]: **y**

Cleaning old directories  
/opt/JNPRhadm

-----  
SBR 8.60.50006 cluster MyCluster{0s,2sm,0m,2d}  
on SunOS 5.10 Generic\_141444-09 node MyNode\_3(d)  
Reviewing Configuration Files...  
-----

/opt/JNPRsbr/radius/install/tmp/dbcluster.rc  
/opt/JNPRsbr/radius/install/tmp/config.ini

```
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
```

```
View (v), accept (a), or reject (r) configuration files: a
```

7. Enter **a** to accept the modified configuration files and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

8. Specify that you want to configure the data node host to autoboot (restart automatically when the operating system is restarted).

#### Enable (e), disable (d), or preserve (p) autoboot scripts [e]: e

A local **/radiusdir/radius/sbrd** script is always created, and **/opt/JNPRhadm/sbrd** is always a symbolic link to this local copy.

- If you enter **e** (enable), the **configure** script copies the local **sbrd** script to **/etc/init.d**, where it is automatically invoked by the OS whenever the OS is stopped or started.
- If you enter **d** (disable), the **configure** script removes all copies of the **sbrd** script from **/etc/init.d**, thus, disabling autoboot for all versions of Steel-Belted Radius Carrier.
- If you enter **p** (preserve), the **configure** script does nothing, thereby leaving your previous autoboot scripts unchanged.

9. Repeat this procedure on each data node in the cluster.
10. Now that the two SBR/management (sm) nodes and two data (d) nodes are configured, start the cluster following the procedure described in [“Initial Cluster Startup Procedure” on page 132](#).

## Initial Cluster Startup Procedure

After all four nodes in the Starter Kit have been installed and configured, we recommend you start the cluster to ensure that everything is working properly, before fully configuring the nodes.

A specific sequence of steps is required to gracefully bring up and shut down the cluster. Do not proceed to the next node until you are sure the process has started properly and the node is connected. The following procedure provides the general steps for starting the cluster in this example. For complete details, see [“When and How to Restart Session State Register Nodes, Hosts, and Clusters” on page 196](#).

1. On the *SBR/management* node that you installed first (**MyNode\_1** in the example), start the **ssr** process:
  - a. Log in as root.
  - b. Change directories to `/opt/JNPRsbr/radius/`.
  - c. Execute:  
`./sbrd start ssr`
  - d. Execute:  
`./sbrd status`
  - e. Examine each line and ensure the SSR process is running without error.
2. Repeat the sequence of commands on the second combined *SBR/management* node (**MyNode\_2** in the example), start the **ssr** process:
  - a. Log in as root.
  - b. Change directories to `/opt/JNPRsbr/radius/`.
  - c. Execute:  
`./sbrd start ssr`
  - d. Execute:  
`./sbrd status`
  - e. Examine each line and ensure the SSR process is running without error.
3. Repeat the sequence of commands on the first *data node* that you installed (**MyNode\_3** in the example), start the **ssr** process:
  - a. Log in as root.
  - b. Change directories to `/opt/JNPRsbr/radius/`.
  - c. Execute:  
`./sbrd start ssr`
  - d. Execute:  
`./sbrd status`
  - e. Examine each line and ensure the SSR process is running without error.
4. Repeat the sequence of commands on the second *data node* that you installed (**MyNode\_4** in the example):
  - a. Log in as root.
  - b. Change directories to `/opt/JNPRsbr/radius/`.

c. Execute:

```
./sbrd start ssr
```

d. Execute:

```
./sbrd status
```

e. Examine each line and ensure the SSR process is running without error.

In the preceding steps, each time the **sbrd status** command is executed, results similar to this example should be displayed:

```
hadmUser$>/opt/JNPRsbr/radius/sbrd status
```

```
[ndbd(NDB)]      2 node(s)
id=1    @172.28.84.163  (mysql-5.7.25  ndb-7.6.9, Nodegroup: 0, Master)
id=2    @172.28.84.113  (mysql-5.7.25  ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)]  2 node(s)
id=51    @172.28.84.36  (mysql-5.7.25  ndb-7.6.9)
id=52    @172.28.84.166  (mysql-5.7.25  ndb-7.6.9)
```

```
[mysqld(API)]    4 node(s)
id=61    (not connected, accepting connect from 172.28.84.36)
id=62    (not connected, accepting connect from 172.28.84.166)
id=100   (not connected, accepting connect from 172.28.84.36)
id=101   (not connected, accepting connect from 172.28.84.166)
```

Examine the lines starting with **id=**, and verify that there are no references to **starting**, **connecting**, or **not connected**. Any of these references indicate the process either has not finished starting, or the node is not connected properly. You may need to execute the **sbrd status** command more than once because it only shows a snapshot of activity; the display does not refresh automatically.

5. Create the database on *every management node*. Management nodes include both normal *management (m) nodes* and combination *SBR/management (sm) nodes*.



This step creates a basic database on each management node in the cluster. Alternatively, you can create a sample database, or customize the database for your particular environment:

- If you want to create a sample database, follow the procedure in [“Testing the Installation with DemoSetup.sh” on page 482.](#)
- If you want to customize the database, see [“Customizing the SSR Database Current Sessions Table” on page 144.](#)

If you choose one of these two options as opposed to performing this step, be sure to come back to this procedure and complete the remaining steps.

**NOTE:** Except when migrating from a temporary cluster, all SSR processes must be *up* on all SSR nodes [sm, m, d] and all SBR processes must be *down* on all SBR nodes [s, sm] in order to execute **CreateDB.sh**.

On each and every management node:

- Log in as hadm.
- Change directories to **/opt/JNPRhadm/**.
- Execute:  
**CreateDB.sh**

6. Configure all server configuration files for your environment.

Complete the configuration of all server initialization (**.ini**) files, authentication (**.aut**) files, accounting (**.acc**) files, as well as configure any proxy setup you may require.

Carefully review the *SBR Carrier Reference Guide* and configure all files for your environment before you start the RADIUS process. Also review the *SBR Carrier Administration and Configuration Guide*, and plan the configuration steps for your particular environment.

You cannot connect to the servers in the cluster with Web GUI until the RADIUS process is started; however, we recommend you plan out the administration of the server before starting the RADIUS process.

See [“Recommendations before Configuring the Cluster” on page 137](#) for general configuration recommendations.

After you have completed the configuration of the various configuration files described in the *SBR Carrier Reference Guide*, remember to come back to this procedure and complete the step in bringing up the cluster.

7. Configure at least one IP address pool and one range using the SSR Administration Scripts. See [“Testing the Installation with DemoSetup.sh” on page 482.](#) Also see the section on *Session State Register Administration* in the *SBR Carrier Administration and Configuration Guide*.

8. Start the RADIUS process on each and every SBR nodes, *one at a time*.

SBR nodes include both *SBR (s) nodes* and *SBR/management nodes (sm)*.

- a. Log in as root to each SBR (s) node and each SBR/management (sm) node.
- b. Change directories to **/opt/JNPRsbr/radius/**.
- c. Execute:  
**./sbrd start radius**
- d. Execute:  
**./sbrd status**

```
[ndbd(NDB)]      2 node(s)
id=1      @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, *)
id=2      @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)

[ndb_mgmd(MGM)]  2 node(s)
id=51     @172.28.84.36   (mysql-5.7.25 ndb-7.6.9)
id=52     @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)

[mysqld(API)]    4 node(s)
id=61     @172.28.84.36   (mysql-5.7.25 ndb-7.6.9)
id=62     @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)
id=100    @172.28.84.36   (mysql-5.7.25 ndb-7.6.9)
id=101    @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)

-----

Current state of network interfaces:

tcp      0      0 0.0.0.0:1812          0.0.0.0:*             LISTEN
tcp      0      0 0.0.0.0:1813          0.0.0.0:*             LISTEN
udp      0      0 172.28.84.36:1645     0.0.0.0:*
udp      0      0 172.28.84.36:1646     0.0.0.0:*
udp      0      0 172.28.84.36:1812     0.0.0.0:*
udp      0      0 172.28.84.36:1813     0.0.0.0:*

-----

hadm      16788 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
hadm      16849 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
hadm      17194 /opt/JNPRmysql/install/bin/mysqld
--basedir=/opt/JNPRmysql/install --datadir=/opt/JNPRmysql/data
--plugin-dir=/opt/JNPRmysql/install/lib/plugin
--log-error=/opt/JNPRmysqld/mysqld_safe.err
--pid-file=/opt/JNPRmysqld/mysqld.pid --socket=/opt/JNPRhadm/.mysql.sock
--port=3001
```

```
root      17683 radius sbr.xml
root      17723 webserver
```

- e. Examine each line and ensure the RADIUS process is running without error.
  - f. Repeat this process until the RADIUS process is started and running without error on each and every SBR node.
9. Complete the configuration of the cluster nodes using Web GUI. See [“Basic SBR Carrier Node Configuration” on page 138](#). For complete details, see the *SBR Carrier Administration and Configuration Guide*.

## Recommendations before Configuring the Cluster

Before proceeding to configure the SBR Carrier nodes and the SSR database, remember that you are configuring a set of machines to act as one entity. You must take some special steps to keep them all working together.

We recommend that you follow these guidelines:

- Always create a backup file before making any modifications to any file.
- Remember that certain files must be the same on all nodes of the same type. To make this easier to manage, include a comment in modified files that records the date and type of modification.
- To implement changes, one or more processes often need to be restarted to force modified files to be read and used. Controlling restarts precisely is easier if you become familiar with the **sbrd** script and its parameters.

# Basic SBR Carrier Node Configuration

## IN THIS CHAPTER

- [Launching Web GUI | 138](#)
- [Configuring the Server | 140](#)
- [Configuring SNMP | 141](#)
- [Setting Up IP Address Pools | 142](#)

This chapter contains basic information about setting up Steel-Belted Radius nodes in a Session State Register. This chapter contains these topics:

## Launching Web GUI

To test whether the server can be accessed by a management workstation, launch the Web GUI.

To launch the Web GUI:

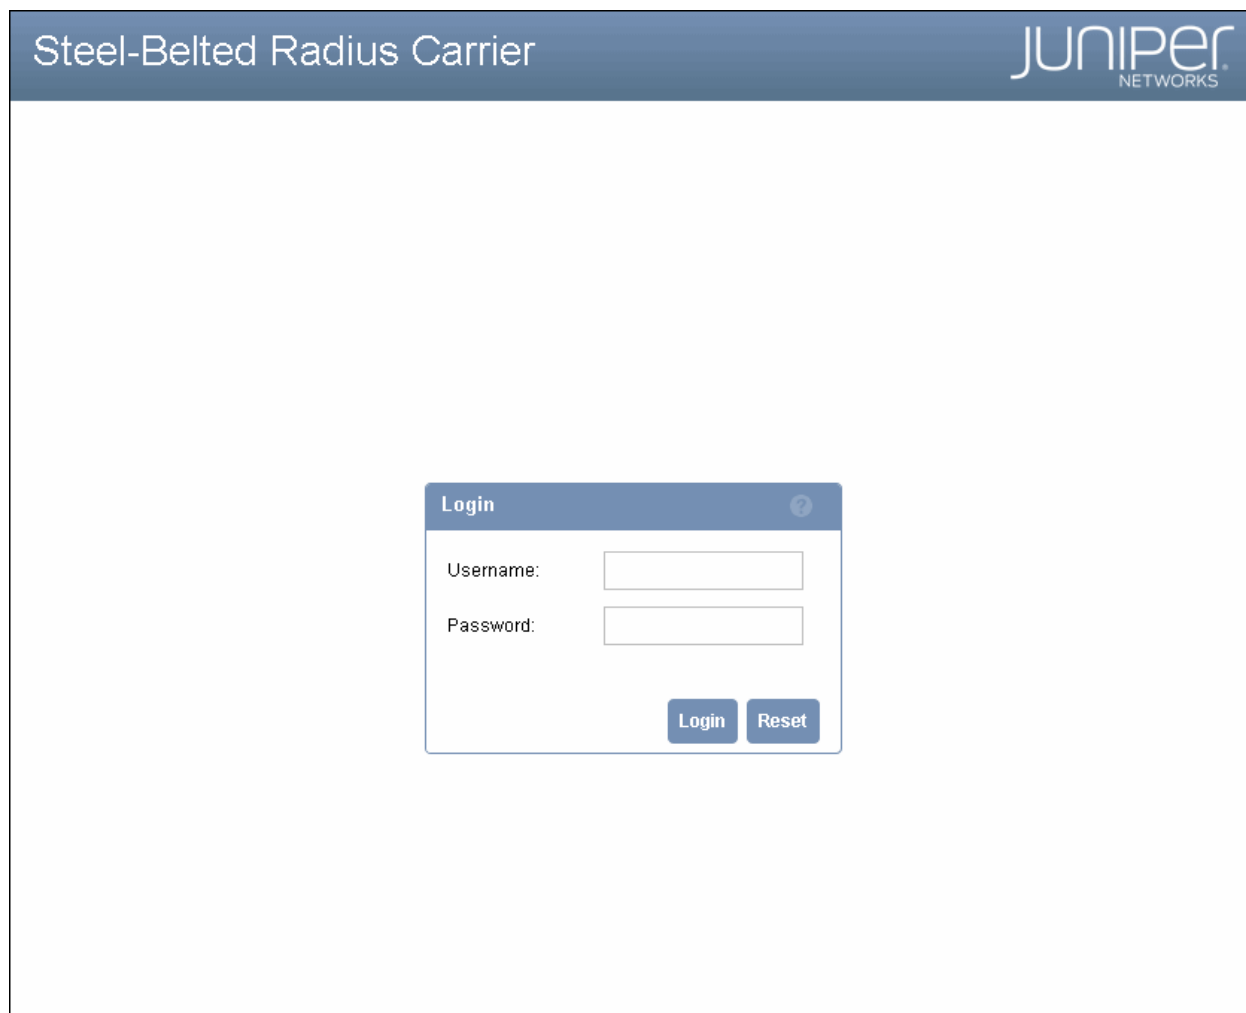
1. Open a browser connection to the SBR Carrier server you want to administer.
  - To administer a SBR Carrier server running on your local host, enter **https://localhost:2909**, where the port assignment of 2909 is the SBR Carrier's default TCP port for administration connections.
  - To administer a SBR Carrier server running on a remote host, enter **https://server:2909**, where the port assignment of 2909 is the SBR Carrier's default TCP port for administration connections. For example:

```
https://192.168.24.15:2909/
```

**NOTE:** Make sure that you access the Web GUI using HTTPS instead of HTTP.

The **Login** page ([Figure 13 on page 139](#)) appears.

Figure 13: Login Page



Steel-Belted Radius Carrier

JUNIPER  
NETWORKS

Login

Username:

Password:

Login Reset

**NOTE:** You can use the **https.port** parameter in the **https.ini** file located in the **/opt/JNPRsbr/radius/website/webserver/jetty/start.d/** path to change the default port.



**CAUTION:** Using a non-default port or modifying other parameters in the **/opt/JNPRsbr/radius/website/webserver/** directory incorrectly could cause communication problems between the Web GUI and the SBR Carrier server.

2. Enter your administrator username in the **Username** field.
3. Enter your login password in the **Password** field.
4. Click **Login**.

When you click **Login**, Web GUI establishes an HTTPS connection with the local or remote server. The Web GUI displays an error message if the connection cannot be established.

**NOTE:** If a timeout occurs, verify that the Web browsers are running on the target server using the **sbrd status** command and that it is listening on the administration port you entered in the URL; that the port is not blocked.

Web GUI verifies that the username you entered exists in the **access.ini** file. If the username is found, Web GUI validates the password you entered against a local or remote password database.

When you connect to a server, the **Home** page lists various features of the running server, such as version, platform on which it is running, IP address, available authentication methods, license information, and any initialization errors that might have occurred.

## Configuring the Server

After you have installed the Steel-Belted Radius Carrier software on the server, have added the appropriate licenses, and can work with Web GUI, you can begin configuring the software. The specific steps you must perform depend on your network's authentication and accounting needs.

The basic steps for configuring the Steel-Belted Radius Carrier environment include:

1. Configure each of your RADIUS client devices to communicate with your Steel-Belted Radius Carrier server. To do this, you must log in to each device and run its configuration interface.
2. Use the **RADIUS Clients List** page in Web GUI to configure the server to communicate with each RADIUS client. Details are in the *Administering RADIUS Clients and Client Groups* section of the *SBR Carrier Administration and Configuration Guide*.
3. If the clients use RADIUS Location Groups or IP Address Pools, each of those entities must exist before the clients are configured. See the sections on *Administering RADIUS Location Groups* and *Administering Address Pools* in the *SBR Carrier Administration and Configuration Guide* for instructions on setting these up.
4. Use the Users panel to identify the users or groups of users who are permitted to access the RADIUS clients.

Specify user attributes by selecting them in the Users panel or by creating user profiles in the Profiles dialog.

For more information, see the *SBR Carrier Administration and Configuration Guide*.

## Configuring SNMP

Steel-Belted Radius Carrier runs its own SNMP agent, but other SNMP agents run on most servers. In general, only one application can use a socket port; they are not shared resources.

During installation, the **configure** script prompts you for SNMP setup information, including an opportunity to specify a port other than the default 161 port that is usually in use by the Solaris SNMP agent.

- If you already provided an alternate port during that setup step, you can skip the following procedure about how to change the port number, but remember to set your MIB browser to listen on the port you specified.
- If you know that other agents already use port 161 but you did not specify an alternate during installation, change the Steel-Belted Radius Carrier port assignment by editing both ***radiusdir*** ***/snmp/conf/jnprsnmpd.conf*** and ***radiusdir*** ***/snmp/bin/testagent.sh***. Remember to check your MIB browser to determine whether it also needs adjustment to communicate with the SBR Carrier server.

To change the port, edit the SBR Carrier SNMP configuration files listed in [Table 19 on page 141](#):

**Table 19: SNMP Configuration Files**

Filename	Function
<b><i>jnprsnmpd.conf</i></b>	Stores settings for the Steel-Belted Radius Carrier SNMP agent.
<b><i>testagent.sh</i></b>	Test script that verifies the Steel-Belted Radius Carrier SNMP agent is operating correctly.

1. Edit *radiusdir /snmp/conf/jnprsnmpd.conf* to change the port number.

The *jnprsnmpd.conf* file is self-documenting. For more information, see the section on *SNMP Configuration Overview* in the *SBR Carrier Reference Guide*.



**CAUTION:** The *jnprsnmpd.conf* file is very sensitive to stray white space and the order in which sections and parameters appear. Mistakes in this file can disable SNMP.

- Make sure to make a backup copy of the file before making any changes.
- While editing the file, do not to make any unnecessary changes. Follow the embedded examples as closely as possible.
- When specifying networks, as in 172.28.68.0/24 in the **com2sec mynetwork 172.28.68.0/24 public** line, the trailing 32-x bits of the IP address must be zero as specified by the trailing /x notation. For example, 32-24=8 bits in this case.

2. Make the same port number change in *radiusdir /snmp/bin/testagent.sh* script, which is used to test the agent.
3. After making the change, restart either the Steel-Belted Radius Carrier server process or just its SNMP daemon.

Execute:

```
/etc/init.d/init.jnprsnmpd start
```

4. If necessary, set up your SNMP browser to listen on the new port.
5. To verify that the *jnprsnmpd* SNMP agent functions, run the *radiusdir /snmp/bin/testagent.sh* script.

**NOTE:** Refer to the *SBR Carrier Administration and Configuration Guide* for more information about configuring the SBR SNMP agent.

## Setting Up IP Address Pools

Set up IP address pools for the SSR database from a cluster management node now.

The same basic concepts discussed in *Administering Address Pools* in the *SBR Carrier Administration and Configuration Guide* apply to address pools maintained in the SSR database; however, the Web GUI is not used to maintain the pools in the SSR database.



In the SSR cluster, you set up and administer IP pools using a set of scripts discussed in the *Session State Register Administration* section of the *SBR Carrier Administration and Configuration Guide*.

**BEST PRACTICE:** We recommend that each SBR Carrier node have its own emergency IP address range, a local pool of addresses that is not shared. If a failover situation takes the shared IP address pool offline, each SBR Carrier node can still function by defining and allocating addresses from its emergency pool. Make sure addresses in the emergency pool are not in the same range of addresses used by the shared SSR IP address pools.

# Customizing the SSR Database Current Sessions Table

## IN THIS CHAPTER

- [Current Sessions Table Overview | 144](#)
- [Sessions in the CST | 153](#)
- [Customizing the CST | 154](#)
- [Current Sessions Table Fields | 156](#)
- [SSR Datatypes | 172](#)

This chapter describes modifying the **CurrentSessions.sql** file that control how the Current Sessions and other tables are configured. It describes fields and what modifications you can make to the files. These topics are in this chapter:

Changes to the SSR database affect all SBR Carrier nodes that are part of the cluster.



**CAUTION:** All SBR Carrier nodes that are part of the cluster must be in management mode before you make changes to the IP pool address database.

If you are working on a cluster that is fully functional, be sure to review the section on *Session State Register Administration* in the *SBR Carrier Administration and Configuration Guide* for information about managing cluster nodes while working on the database.

## Current Sessions Table Overview

The shared memory engine of the data cluster contains the current sessions table (CST), which lists current sessions. The CST is the largest and most active element of the database. The CST contains one record, or row, for each session created or updated by any of the Steel-Belted Radius Carrier servers using the data cluster.

The types of data that a row contains are controlled by adding or deleting fields (which form columns in the CST) to the `/opt/JNPRhadm/CurrentSessions.sql` file. Adding a column creates a new field for each record in the table. Within the configuration file, columns are defined as belonging to specific types.

The information stored in the CST is customizable. To ensure that you can identify sessions accurately and format CoA/DM requests with the appropriate attributes for the services you are supporting in your network, we recommend that you customize the information stored in the CST for your particular network environment. For example, before you retrieve a value from the CST for use in an attribute in a CoA message, you must first map the CST field to a RADIUS attribute in the `dbc_mapping.xml` file. In the following example, the **FunkOuterUserName** field is mapped to the **Original-User-Name** attribute using the **attribute** element to enable retrieving the value in the **FunkOuterUserName** field as a result of a session control request. The **queryAttribute** child element is used for indexing the CST while querying sessions using the **SessionControl.sh** script or Web GUI. If you want to query sessions using the attribute, then you need to specify the attribute in the **queryAttribute** child element.

```
<attributeMapping field="FunkOuterUserName" attribute="Original-User-Name">
  <queryAttribute name="Original-User-Name"/>
</attributeMapping>
```

When Steel-Belted Radius Carrier SSR is installed, the default CST schema contains Core, Required, and Optional System fields. By default, the System Core fields and the **Sbr\_UserName** System Optional field are active, with other System fields commented out. When **CreateDB.sh** is invoked, it includes only active fields in the CST.

If you modify the file, remember:

- All System Core fields must be active.
- System Feature fields must be active if Steel-Belted Radius Carrier is configured to support their corresponding features. To activate an inactive System Feature field, you must uncomment the entry for the field in the **CurrentSessions.sql** file. (If the System Feature field required for a feature is inactive, the feature is disabled and an error message is written to the RADIUS log file.)
- You can comment out the **Sbr\_UserName** field if it is not required.
- After **CurrentSessions.sql** is modified:
  - a. The file must be distributed to all management nodes in the cluster.
  - b. The current running database must be removed. As hadm, run **DestroyDB.sh** on one management node.
  - c. The new database must be created. As hadm, run **CreateDB.sh** on all management nodes to rebuild the database.

**NOTE:** You can change the order of fields, except the System Core fields, in the **CurrentSessions.sql** file. We strongly recommend that you not edit the System Core fields as this may result in unexpected session deletions.

[Table 20 on page 147](#) describes each section and each section's fields:

Table 20: Default Current Sessions Table Fields

Field Type	CurrentSessions.sql Entry
<p>System Core fields are mandatory fields that must be present in every CST.</p> <p>See <a href="#">Table 21 on page 158</a> for details on each field.</p>	

Table 20: Default Current Sessions Table Fields (continued)

Field Type	CurrentSessions.sql Entry
	<pre> #----- ##### This CST schema is customizable, but CONSULT THE DOCUMENTATION FIRST!! # #####  CREATE TABLE Sbr_CurrentSessions (  #-----          SYSTEM CORE FIELDS  #-----          Sbr_UniqueSessionId      BINARY(16)                                 NOT NULL /*UNIQUE*/,  *is* "UNIQUE", but dont declare it so (MySQL bug)         Sbr_CreationTime         TIMESTAMP                                 NOT NULL DEFAULT 0, Sbr_ExpirationTime             TIMESTAMP                                 NOT NULL DEFAULT 0,         Sbr_Ipv4Address           INT UNSIGNED                                 DEFAULT NULL, Sbr_IpPoolOrdinal              SMALLINT UNSIGNED                                 DEFAULT NULL,         Sbr_NasName               VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci                                 NOT NULL,         Sbr_SessionState          TINYINT UNSIGNED                                 NOT NULL DEFAULT 0,         Sbr_UserConcurrencyId     VARCHAR(84) CHARSET utf8 COLLATE utf8_general_ci                                 DEFAULT NULL,          Sbr_MobileIpType          TINYINT UNSIGNED                                 DEFAULT 0, Sbr_3gpp2ReqType               INT UNSIGNED                                 DEFAULT 0,         Sbr_WimaxClientType       TINYINT UNSIGNED                                 NOT NULL DEFAULT 0, </pre>

Table 20: Default Current Sessions Table Fields (continued)

Field Type	CurrentSessions.sql Entry
	<pre> Sbr_WimaxAcctFlows      VARBINARY(4095)                         DEFAULT NULL, Sbr_3gpp2HomeAgentAddr  INT UNSIGNED                         DEFAULT NULL, #-----  #-----  # SYSTEM CORE FIELD added for RFC6911 IPv6Address at index 31 #-----  Sbr_Ipv6Address          VARBINARY(2047)                         DEFAULT NULL, </pre>
<p>System Feature fields are static fields that must be present if, and only if, specific Steel-Belted Radius Carrier features are enabled.</p> <p>See <a href="#">Table 22 on page 160</a> for details on each field.</p>	<pre> SYSTEM FEATURE FIELDS #----- #Sbr_AcctFileCarryover  VARBINARY(1024)                         DEFAULT NULL, Sbr_AcctAutoStop        VARBINARY(1023)                         DEFAULT NULL, Sbr_SessionTimeout      INT UNSIGNED                         DEFAULT NULL, Sbr_ClassAttribute      VARBINARY(1024)                         DEFAULT NULL, #Sbr_UniqueSessionId_Hex CHAR(32) CHARSET ascii COLLATE ascii_general_ci                         DEFAULT NULL, INDEX Sbr_UniqueSessionId_Hex_Idx USING HASH (Sbr_UniqueSessionId_Hex), #----- </pre>

Table 20: Default Current Sessions Table Fields (continued)

Field Type	CurrentSessions.sql Entry
<p>System Optional fields are static fields that you can use or comment out. By default, all fields are enabled.</p> <p>See <a href="#">Table 23 on page 161</a> for details on each field.</p>	<pre>SYSTEM OPTIONAL FIELDS #----- Sbr_UserName          VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL, Sbr_AcctSessionId     VARCHAR(48) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL, Sbr_TransactionId     BINARY(12) DEFAULT NULL, Sbr_NasPortType       INT UNSIGNED DEFAULT NULL, Sbr_NasPort           INT UNSIGNED DEFAULT NULL, Sbr_CallingStationId  VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL, Sbr_CalledStationId   VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL, Sbr_MobileCorrelationId VARCHAR(32) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL, #-----  #-----  # SYSTEM OPTIONAL FIELDS added for Ipv6Prefix at index 30 #-----  Sbr_Ipv6Prefix        VARBINARY(16) DEFAULT NULL,</pre>



Table 20: Default Current Sessions Table Fields (continued)

Field Type	CurrentSessions.sql Entry
<p>RadAttr fields are dynamic custom fields that you can use to capture raw information in incoming or outgoing RADIUS attributes. The content of each field in the CST comes from SBR Carrier attributes.</p> <p>You can modify the three template fields if you need to add custom fields.</p> <p>See <a href="#">Table 24 on page 163</a> for details on each field.</p>	<pre> ADMIN RADATTR FIELDS #----- WimaxSessionId          VARBINARY(32) DEFAULT NULL, AcctMultiSessionId      VARCHAR(32) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL, FunkOuterUserName       VARCHAR(84) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL, #RadAttrField1          INT UNSIGNED                         DEFAULT NULL, #RadAttrField2          VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci                         DEFAULT NULL, #RadAttrField3          VARBINARY(32)                         DEFAULT NULL, #----- </pre>
<p>Private fields are dynamic custom fields that you can use to capture information of any kind from third-party applications. SBR Carrier does not populate private fields, so you can modify the template fields if you need to add private fields.</p>	<pre> ADMIN PRIVATE FIELDS #----- #PrivateField1          INT UNSIGNED                         DEFAULT NULL, #PrivateField2          VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci                         DEFAULT NULL, #PrivateField3          BINARY(32)                         DEFAULT NULL, #----- </pre>

Table 20: Default Current Sessions Table Fields (continued)

Field Type	CurrentSessions.sql Entry
<p>System Keys and Indexes of columns in the table list tools that are required for Session State Register admin scripts and shell functions. They may also be used by other applications making calls through the NDBAPI or SQL queries.</p> <ul style="list-style-type: none"> <li>• A single-column index may be used by SBR Carrier front ends, SSR itself, or the LDAP Configuration Interface.</li> <li>• A multi-column index may only be used within the Session State Register cluster or from a third party application. It may not be used by SBR Carrier front ends, SSR itself, or the LDAP Configuration Interface.</li> </ul> <p>Because of these limitations, you cannot use a multi-column index for the session lookup in a CoA/DM request (for example). However, you can add a private field, index with just that one field, and use that index in your CoA/DM query strings.</p>	<pre> SYSTEM KEYS/INDEXES #----- PRIMARY KEY USING HASH                (Sbr_UniqueSessionId), INDEX Sbr_NasName_Idx                (Sbr_NasName), INDEX Sbr_ExpirationTime_Idx        (Sbr_ExpirationTime), INDEX AcctMultiSessionId_Idx        (AcctMultiSessionId), INDEX Sbr_MobileCorrelationId_Idx   (Sbr_MobileCorrelationId), INDEX Sbr_AcctSessionId_Idx         (Sbr_AcctSessionId), INDEX Sbr_UserName_Idx              (Sbr_UserName), INDEX Sbr_CallingStationId_Idx      (Sbr_CallingStationId), INDEX Sbr_CalledStationId_Idx       (Sbr_CalledStationId), INDEX Sbr_Ipv4Address_Idx           (Sbr_Ipv4Address), INDEX Sbr_TransactionId_Idx         (Sbr_TransactionId) #----- </pre>

Table 20: Default Current Sessions Table Fields (*continued*)

Field Type	CurrentSessions.sql Entry
Admin keys and indexes must be present in every CST.	<pre> ADMIN KEYS/INDEXES #----- #,     INDEX Some_Idx USING HASH (What, Ever) ) ENGINE = ndbcluster  NOTE: CreateDB.sh fiddles with this line! ;  #===== </pre>

## Sessions in the CST

The CST maintains a record for most, but not all, sessions.

- When only System fields are active in the CST, Steel-Belted Radius Carrier creates a CST record using a two-stage process:
  - a. At authentication, when a RADIUS authentication request is processed successfully and a RADIUS authentication response is sent. Tracked resources are allocated for the session. These sessions are phantom sessions, and are typically promoted to a real session when a RADIUS Accounting-Start message arrives.
  - b. At accounting start, when a RADIUS Accounting-Start message for the session is received.
- If RadAttr fields are active in the CST, Steel-Belted Radius Carrier creates a phantom session for every authentication request (not just those containing attributes to be captured in a RadAttr field). Steel-Belted Radius Carrier then records RADIUS attributes, including the RadAttr fields, in the CST.

## Acknowledging Authentication and Accounting Requests when the CST Cannot Be Contacted

By default, accounting requests are acknowledged even if the session database cannot be contacted. To cause accounting requests to be discarded when the session database cannot be contacted, which may be desirable when using load-balancing equipment, modify **radius.ini** as follows:

```

[Configuration] Section
DiscardAccountingRequestOnCstFailure = 1

```

- If set to 1, accounting requests (start, stop, on, off, and interim) are discarded when the session database cannot be contacted.

- If set to 0, accounting requests (start, stop, on, off, and interim) are acknowledged when the session database cannot be contacted.

Similarly, to cause the discard of authentication requests that contact the session database to assign resources (such as IP address assignment or concurrency), modify **radius.ini** as follows:

```
[Configuration] Section
DiscardAccessRequestOnCstFailure = 1
```

- If set to 1, authentication requests requiring access to the session database are discarded when the session database cannot be contacted.
- If set to 0, SBR Carrier sends an Access-Reject when the session database cannot be contacted.

**NOTE:** Operation is not affected for requests not requiring session database access.

## Current Sessions Table Display

You can display the working CST by executing the **ShowSessions.sh** script. The format used by the **ShowSessions.sh** tool is specified in the SQL data-display language. The **ShowSessions.sh** script may be modified.

When Steel-Belted Radius Carrier SSR is installed, the **ShowSessions.sh** tool is configured to work with the default CST schema. The **ShowSessions.sh** tool contains display commands for all System fields, which are commented-in or commented-out appropriately. If you comment or uncomment entries in the SQL schema to meet site requirements, you must comment or uncomment corresponding entries in **ShowSessions.sh** too.

[Figure 14 on page 175](#) shows sample **ShowSessions.sh** output.

## Customizing the CST

A key benefit of Session State Register is that you can modify the CST to meet a site's requirements by adding or modifying session fields. However, if you make too many modifications, or make them incorrectly, problems can occur. Review the following notes before making any modifications, and test your changes as you go.

In general, try to use as many defaults as possible, and add only what is required.

## Propagating a Changed CurrentSessions.sql File

Because you are working in a cluster environment where multiple machines use the same data, making changes to the CST table is not a trivial task. Each time you change the configuration file, it must be distributed to all management nodes. Then you must eliminate the existing database (with the **DestroyDB.sh** command) and create a new database. (As hadm, execute the **CreateDB.sh** command on all management nodes.)

## Performance and Capacity Considerations

Each field that you add to the CST increases the size of every session record stored in the CST. Larger tables increase the amount of memory required by the database. As a result, adding fields to the CST limits the number of simultaneously active sessions that can be stored in the CST, given fixed memory and disk space. Adding fields to the CST slows wire communications and internal processing time.

The Steel-Belted Radius Carrier SSR implementation limits the maximum supported size of a CST field to 4096 bytes. This limit matches the maximum size of a RADIUS packet as defined in RFC 2865. RFC 2865 also limits each packet's individual RADIUS attributes to a maximum of 253 bytes.

In practice, RADIUS attributes and packets are usually much smaller than 4096 bytes. The entire cluster runs more efficiently when fields are tuned to the size of a site's RADIUS attributes and corresponding RadAttr fields.

## Additional Keys

The default **CurrentSessions.sql** contains the primary key for the CST and other keys (indexes) required for Steel-Belted Radius Carrier operation. You can comment unused keys out, but we recommend that you do not delete any of these in case they need to be restored. You can add additional keys on any fields to support faster lookups for site-specific tools or third-party products. Each additional key adds overhead every time a row (session) is added to or removed from the CST.

## Stored Procedures

SSR includes basic stored procedures defined in the **StoredRoutines.sql** file. You cannot modify or delete these stored procedures, but you can define new stored procedures to perform special handling of RadAttr fields for tools you may use or create. Creation of custom stored procedures requires advanced knowledge of SQL.

Remember that SQL queries tend to run more slowly, and to require more resources, than queries and scripts (such as the admin scripts) that use NDBAPI.

## Customized CST Applications

You can write custom SQL-based tools to query the SSR database. You can base these on the SSR database CLI (just as **ShowSessions.sh** is), or the tools can talk directly to a SSR management server mysqld daemon. You can also define new stored procedures in **StoredRoutines.sql** and may create new NDBAPI applications.

## Current Sessions Table Fields

The CST and the **CurrentSessions.sql** contain several types of fields:

- “System Core Fields” on page 156
- “System Feature Fields” on page 159
- “System Optional Fields” on page 160
- “RadAttr Fields” on page 162
- “Admin Private Fields” on page 170

The table also contains settings for keys and indexes:

- “System Keys and Indexes” on page 170
- “Admin Keys and Indexes” on page 171

### System Core Fields

The System Core fields are required for basic operation of SBR Carrier SSR.



**CAUTION:** We strongly recommend that you not edit this section of the file. If any required Core field is not present in the CST, Steel-Belted Radius Carrier SSR cannot boot.

The System Core Fields section of **CurrentSessions.sql** looks like this:

```
#=====
```

```
#####
This CST schema is customizable, but CONSULT THE DOCUMENTATION FIRST!! #
#####
```

```
CREATE TABLE Sbr_CurrentSessions
(
  #-----
  SYSTEM CORE FIELDS
  #-----
  Sbr_UniqueSessionId      BINARY(16)
                           NOT NULL /*UNIQUE*/,  *is* "UNIQUE", but dont
declare it so (MySQL bug)
  Sbr_CreationTime         TIMESTAMP
                           NOT NULL DEFAULT 0,
  Sbr_ExpirationTime       TIMESTAMP
                           NOT NULL DEFAULT 0,
  Sbr_Ipv4Address          INT UNSIGNED
                           DEFAULT NULL,
  Sbr_IpPoolOrdinal        SMALLINT UNSIGNED
                           DEFAULT NULL,
  Sbr_NasName              VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci
                           NOT NULL,
  Sbr_SessionState         TINYINT UNSIGNED
                           NOT NULL DEFAULT 0,
  Sbr_UserConcurrencyId    VARCHAR(84) CHARSET utf8 COLLATE utf8_general_ci
                           DEFAULT NULL,
  Sbr_MobileIpType         TINYINT UNSIGNED
                           DEFAULT 0,
  Sbr_3gpp2ReqType         INT UNSIGNED
                           DEFAULT 0,
  Sbr_WimaxClientType      TINYINT UNSIGNED
                           NOT NULL DEFAULT 0,
  Sbr_WimaxAcctFlows       VARBINARY(4095)
                           DEFAULT NULL,
  Sbr_3gpp2HomeAgentAddr   INT UNSIGNED
                           DEFAULT NULL,
  #-----
  # SYSTEM CORE FIELD added for RFC6911 IPv6Address at index 31
  #-----
```

Sbr_Ipv6Address	VARBINARY( 2047 ) DEFAULT NULL ,
-----------------	-------------------------------------

The datatypes and sizes for System Core fields are specified [Table 21 on page 158](#). Some of these datatypes and sizes are fixed. The ones that may be changed have default values specified in the table. These defaults are a starting point to use during pre-production prototyping.

**Table 21: System Core Fields and Specifications**

Name	Description	SQL Datatype and Size
<b>Sbr_UniqueSessionId</b>	Internally generated identifier unique within the scope of the CST. The value is generated by a proprietary algorithm. The field acts at the Primary Key for the CST.	BINARY(16).  Fixed.
<b>Sbr_CreationTime</b>	The time at which this session was created.	TIMESTAMP.  Fixed.
<b>Sbr_ExpirationTime</b>	The time at which this session expires. The value may be updated during operation, for example, by session keepalive messages.	TIMESTAMP.  Fixed.
<b>Sbr_Ipv4Address</b>	The IPv4 address associated with this session (if any). Compare the RADIUS attribute Framed-IP-Address.	INT UNSIGNED.  Fixed (interpreted as IPv4 address).
<b>Sbr_IpPoolOrdinal</b>	Internal implementation artifact, used by Steel-Belted Radius Carrier SSR IPv4 address pools.	SMALLINT UNSIGNED.  Fixed (determined by IpPools table).



Table 21: System Core Fields and Specifications (continued)

Name	Description	SQL Datatype and Size
<b>Sbr_NasName</b>	<p>NAS name as assigned by Steel-Belted Radius Carrier, using (in preferential order):</p> <ol style="list-style-type: none"> <li>1. RADIUS attribute NAS-IP-Address (or NAS-IPv6-Address, if available)</li> <li>2. RADIUS attribute NAS-Identifier</li> <li>3. Actual IPv4 (or IPv6, if available) address from which the TCP/IP packet was sent.</li> </ol>	<p>VARCHAR(24) utf8.</p> <p>Customizable type and size.</p>
<b>Sbr_SessionState</b>	Internal implementation artifact, used to manage status of sessions.	<p>TINYINT UNSIGNED.</p> <p>Fixed.</p>
<b>Sbr_UserConcurrencyID</b>	Internal implementation artifact (username concatenated with integer), used by Steel-Belted Radius Carrier SSR to manage 'user concurrency' feature.	<p>VARCHAR(32) utf8.</p> <p>Customizable type and size, but must be the same type and size as the Id field in the Sbr_User Concurrency table.</p>
<b>Sbr_Ipv6Address</b>	The IPv6 address associated with this session (if any). Compare the RADIUS attribute Framed-IPv6-Address.	<p>VARBINARY(2047).</p> <p>Fixed (interpreted as IPv6 address).</p>

## System Feature Fields

One or more of the System Feature fields listed in [Table 22 on page 160](#) must be present in the CST for Steel-Belted Radius Carrier to run in a special feature enabled mode. If a required System Feature field is not present in the CST, Steel-Belted Radius Carrier boots, but the special feature is disabled.

The System Feature Fields section of **CurrentSessions.sql** looks like this:

```
DEFAULT NULL,
```

#-----	
SYSTEM FEATURE FIELDS	
#-----	
Sbr_AcctAutoStop	VARBINARY(1023) DEFAULT NULL,
Sbr_SessionTimeout	INT UNSIGNED DEFAULT NULL,
Sbr_ClassAttribute	VARBINARY(1024) DEFAULT NULL,

Table 22 on page 160 describes the System Feature Fields, SQL datatypes, and sizes.

Table 22: System Feature Fields and Specifications

Name	Description	SQL Datatype & Size
Sbr_AcctAutoStop	"Account Auto Stop" feature.	BINARY(1023).  Fixed.
Sbr_SessionTimeout	"Session Timeout on Missed Account Stop" feature. RADIUS attribute Session-Timeout.	INT UNSIGNED.  Fixed.
Sbr_ClassAttribute	"Use Single Class Attribute" feature.	VARBINARY(1024).  Customizable type and size.

System Optional Fields

System optional fields can provide useful information about a session, but none are required, so you may comment out unnecessary fields.

The System Optional Fields section of **CurrentSessions.sql** looks like this:

#-----	
SYSTEM OPTIONAL FIELDS	
#-----	

Sbr_UserName	VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL,
Sbr_AcctSessionId	VARCHAR(48) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL,
Sbr_TransactionId	BINARY(12) DEFAULT NULL,
Sbr_NasPortType	INT UNSIGNED DEFAULT NULL,
Sbr_NasPort	INT UNSIGNED DEFAULT NULL,
Sbr_CallingStationId	VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL,
Sbr_CalledStationId	VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL,
Sbr_MobileCorrelationId	VARCHAR(32) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL,
#-----	
# SYSTEM OPTIONAL FIELDS added for Ipv6Prefix at index 30	
#-----	
Sbr_Ipv6Prefix	VARBINARY(16) DEFAULT NULL,

Table 23 on page 161 describes System Optional Fields, SQL datatypes, and sizes.

Table 23: System Optional Fields and Specifications

Name	Description	SQL Datatype & Size
Sbr_UserName	RADIUS attribute User-Name (as presented by Steel-Belted Radius Carrier).	VARCHAR(24) utf8.  Customizable type and size.
Sbr_TransactionId	Internal, proprietary implementation artifact.	BINARY(12).  Fixed.
Sbr_NasPortType	RADIUS attribute NAS-Port-Type.	INT UNSIGNED.  Fixed.
Sbr_NasPort	RADIUS attribute NAS-Port.	INT UNSIGNED.  Fixed.

Table 23: System Optional Fields and Specifications *(continued)*

Name	Description	SQL Datatype & Size
<b>Sbr_CallingStationId</b>	RADIUS attribute Calling-Station-Id.	VARCHAR(24) utf8.  Customizable type and size.
<b>Sbr_CalledStationId</b>	RADIUS attribute Called-Station-Id.	VARCHAR(24) utf8.  Customizable type and size.
<b>Sbr_AcctSessionId</b>	RADIUS attribute Acct-Session-Id. This attribute is enabled by default.*	VARCHAR(48) utf8.  Customizable type and size.
<b>Sbr_MobileCorrelationId</b>	RADIUS attribute Mobile-Session-Id.	VARCHAR(24) utf8.  Customizable type and size.
<b>Sbr_Ipv6Prefix</b>	The IPv6 prefix associated with this session (if any).	VARBINARY(16)  Fixed

\* The Acct-Session-Id RADIUS attribute is commonly used in CoA/DM packing lists for router requirements. Packing lists represent the router's requirements for a valid DM or CoA request. Attributes that are not in the CST must come from the actual query to fulfill the request.

To have data available for queries using the command line utility (**SessionControl.sh**) for CoA/DM requests, you must configure this attribute in the packing list. If the Sbr\_AcctSessionId does not exist in the CST or is not found by the query itself, the request fails.

## RadAttr Fields

The System Core, System Feature, and System Optional fields are sufficient for most deployments, but if a site has requirements for CST information that cannot be supported by the System fields, you can add RadAttr fields to the CST. RadAttr fields are designed to capture data from RADIUS packet attributes. You must be knowledgeable about RADIUS attributes, how they are used in the network, and which attributes to incorporate into the CST.

You can add as many as 64 RadAttr fields to the CST, but fields beyond the 64th are not populated by Steel-Belted Radius Carrier.

The Admin Radattr Fields section of **CurrentSessions.sql** looks like this:

```
ADMIN RADATTR FIELDS
#-----
WimaxSessionId      VARBINARY(32)
                     DEFAULT NULL,
AcctMultiSessionId   VARCHAR(32) CHARSET utf8 COLLATE utf8_general_ci
                     DEFAULT NULL,
FunkOuterUserName    VARCHAR(84) CHARSET utf8 COLLATE utf8_general_ci
                     DEFAULT NULL,
#RadAttrField1       INT UNSIGNED
                     DEFAULT NULL,
#RadAttrField2       VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci
                     DEFAULT NULL,
#RadAttrField3       VARBINARY(32)
                     DEFAULT NULL,
#-----
```

Table 24 on page 163 describes Admin RadAttr Fields, SQL datatypes, and sizes.

**Table 24: RadAttr Fields and Specifications**

Name	Description	SQL Datatype & Size
WimaxSessionId	Unique WiMAX session ID	VARBINARY(32) Null
AcctMultiSessionId	Unique Accounting Multiple Session ID	VARCHAR(32) CHARSET utf8 COLLATE utf8_general_ci Null
FunkOuterUserName	Unique outer authentication username	VARCHAR(84) CHARSET utf8 COLLATE utf8_general_ci Null
#RadAttrField1	-	INT UNSIGNED Null

Table 24: RadAttr Fields and Specifications (continued)

Name	Description	SQL Datatype & Size
#RadAttrField2	–	VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci  Null
#RadAttrField3	–	VARBINARY(32)  Null

RadAttr fields are similar to some of the System Core and Optional fields, but differ in that the System fields are supported *natively* and are processed by Steel-Belted Radius Carrier. The RadAttr fields work with raw values extracted from RADIUS packets.

When you create a field, follow these guidelines:

- The *RadAttrFieldName* for a field may be any legal SQL name, except that its first three characters cannot be the reserved System prefix Sbr (in any letter case combination).
- Each *RadAttrFieldName* can occur only once in the **sessionTable.ini** file. A *RadAttrFieldName* specifies a column name in the CST, and duplicate column names are illegal in SQL tables.
- You can declare a *RadAttrFieldName* in the CST schema but not declare it in the **sessionTable.ini** file. For example, you can declare the **PrivateFieldName** name in the CST schema, but not in the **sessionTable.ini** because you can define Private fields in the CST to hold any non-SBR Carrier data you want. This is useful where a third-party application accesses the CST.
- Valid **Radius-Attribute-Names** occur in Steel-Belted Radius Carrier dictionary files, as embodied in the SBR Carrier internal primary dictionary (which includes attributes defined in files having **.dct**, **.dcm**, **.dci**, and **.dcx** extensions in the SBR Carrier installation root directory). For information about SBR Carrier dictionaries, refer to the **readme.dct** and **dictiona.dcm** files.



**CAUTION:** You must correctly declare the CST schema in **CurrentSessions.sql** using the MySQL implementation of SQL, or the table may not work correctly. Always check CST modifications in a test environment before putting them into production use.

### **RadAttr Fields and the sessionTable.ini File**

The procedure for adding RadAttr fields is more involved than just editing the **CurrentSessions.sql** file because an additional file, the **sessionTable.ini**, maps RadAttr CST field names to RADIUS attributes. Each addition of a RadAttr field in the **CurrentSessions.sql** file requires a corresponding entry in **sessionTable.ini** that defines the data used to populate the field and when to extract the data from the RADIUS packet.

RADIUS attributes can be generated by Steel-Belted Radius Carrier at four *capture points*, events that apply only to RadAttr fields:

- [AuthRequest]
- [AuthResponse]
- [AcctRequest]
- [AcctResponse]

Each RadAttr CST field requires at least one capture point to identify when RADIUS attributes are loaded in the field. Specify one or more capture-points by placing the field entry in one or more of the four capture point sections of the **sessionTable.ini** file.

One **Radius-Attribute-Name** can be captured multiple times, at different capture-points, by using multiple (different) **RadAttrFieldNames**, all of which are mapped to the same **Radius-Attribute-Name**.

If a RADIUS attribute is not present when a session's RADIUS packet is scanned at a capture point event, the corresponding CST field is not updated. If the specified attribute does not occur in a packet at any capture point, the session displays a null value in the CST field.

If a session experiences multiple occurrences of packets at the same capture-point carrying the same **Radius-Attribute-Name**, the values of subsequent occurrences overwrite the values of preceding occurrences in the **RadAttrFieldName** field. For example, if a session receives multiple Accounting-Start packets, then the AcctRequest and AcctResponse capture-points are exercised more than once for that session, with the result that the captured **Radius-Attribute-Name/ RadAddrFieldnames** are updated multiple times. Thus, the CST always reflects the most recent values from each capture-point.

### **Creating RadAttr Field-to-Attribute Mappings**

To support customizable CST fields requires a new configuration file, (**sessionTable.ini**), which is located in the Steel-Belted Radius Carrier installation root directory. The **sessionTable.ini** file maps **RadAttrFieldName** RadAttr CST fields to **Radius-Attribute-Name** RADIUS attributes:

**RadAttrFieldName = Radius-Attribute-Name**

### **How to Define RadAttr Fields**

Defining a RadAttr field is a three-step procedure:

1. Create the field in the **CurrentSessions.sql** file.

Add the RadAttrFieldName field to the CST schema in **CurrentSessions.sql**.

For example:

```
ADMIN RADATTR FIELDS
#-----
WimaxSessionId          VARBINARY(32)
```

```

DEFAULT NULL,
....
threegppImsi VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci
DEFAULT NULL,
threegppGgsnAddress INT UNSIGNED
DEFAULT NULL,

```

## 2. Edit the **sessionTable.ini** file.

Add the **RadAttrFieldName = Radius-Attribute-Name** field-to-attribute mapping line to the **sessionTable.ini** file in the appropriate capture point section or in several sections.

This example supports the fields you entered Step 1:

```

[AcctRequest]
;=====

```

```

threegppImsi=3GPP-IMSI
threegppGgsnAddress=3GPP-GGSN-Address

```

## 3. Optionally, edit the **ShowSessions.sh** file.

Add the **RadAttrFieldName** to **ShowSessions.sh**, if you want it to appear in the **ShowSessions.sh** command. You can customize the formatting of RadAttr fields, following the guidelines of the System fields in **ShowSessions.sh**.

**NOTE:** Be particularly careful with a RadAttr fields of type **TIMESTAMP**. If one is in the **CST** schema but you fail to declare the field type in **sessionTable.ini**, then no sessions are created in the **CST**. Every purported **CST** record that Steel-Belted Radius Carrier SSR presents to **NDB** has a **NULL** value in the declared **TIMESTAMP** field. Because **TIMESTAMP** fields must be non-**NULL** unless they are explicitly declared as **TIMESTAMP NULL**, **NDB** refuses the request.

We recommend that every RadAttr **TIMESTAMP** field always be declared as **TIMESTAMP NULL** because you can rarely guarantee that the required **RADIUS** attribute will always be present.



After completing the preceding procedure for RadAttr fields and customizing System fields as needed, you must delete the existing database and create a new one that invokes your updated files. Then run **ShowSessions.sh** to verify that everything is working correctly.

### **Multi-Valued Attributes**

The **sessionTable.ini** file supports a special syntax to support multi-valued attributes. The RADIUS protocol permits certain attributes to be *multi-valued*, meaning that two or more instances of the attribute can occur in a single RADIUS packet. System fields in the CST always involve single-valued attributes. RadAttr fields in the CST can be single-valued or multi-valued.

For RadAttr fields in the CST, we support multi-valued attributes with the six MVA facilities defined in the list that follows. Any field-to-attribute mapping (in **sessionTable.ini**) can have its **Radius-Attribute-Name** optionally appended with one (but no more than one) multi-valued attribute @ extension.

#### **MVA-Count**

The MVA-Count extension stores the count (0, 1, or more) of attributes of the specified **Radius-Attribute-Name** (occurring in the RADIUS packet being scanned at the specified capture-point). Declare **RadAttrFieldName** as an INT UNSIGNED field in the CST schema. If you declare it to be [VAR]CHAR or [VAR]BINARY, the field is automatically formatted.

Syntax:

```
RadAttrFieldName = Radius-Attribute-Name@#
```

#### **MVA-Index**

The MVA-Index extension stores the value of the Nth attribute (where N is equal to or greater than 1) of the **Radius-Attribute-Name** occurring in the RADIUS packet being scanned at the specified capture-point. For example, @1 means to store the first value of the specified **Radius-Attribute-Name**. If the packet carries fewer than N attributes of the designated type, a value of null is stored.

In this case, N is a 1-based, not 0-based, integer. Thus, N denotes a positive number in the range 1–2038, expressed in standard base-10 notation. Specifying a value for N outside this range is an error.

Syntax:

```
RadAttrFieldName = Radius-Attribute-Name@N
```

#### **MVA-First**

The MVA-First extension stores the first value of the **Radius-Attribute-Name** attribute occurring in the RADIUS packet being scanned at the specified capture-point. MVA-First extension is synonymous with the **Radius-Attribute-Name**@1 extension.

Syntax:

**RadAttrFieldName** = *Radius-Attribute-Name*@^

#### **MVA-Last**

The MVA-Last extension stores the last value of the **Radius-Attribute-Name** attribute occurring in the RADIUS packet being scanned at the specified capture-point.

Syntax:

**RadAttrFieldName** = *Radius-Attribute-Name*@\$

#### **MVA-CharString-Concat**

The MVA-CharString-Concat extension treats the potentially multiple values of **Radius-Attribute-Name** as text strings. It concatenates them using the specified string "...", as a field delimiter between successive values. Ensure that values of the **Radius-Attribute-Name** attribute are the natural CHAR datatype. If they are not, the appropriate unnatural conversion is applied.

**NOTE:** The **RadAttrFieldName** field must be declared in the CST schema to be of [VAR]CHAR type. If the **RadAttrFieldName** field is of a type other than [VAR]CHAR, Steel-Belted Radius Carrier reports an error.

The value of the delimiter string cannot contain the 2-character sequence "@" (at-sign/double-quote) or internal NULL characters. The delimiter string can consist of any arbitrary (single-byte) characters, such as comma, colon, space, tab, and new line characters. This facility supports C/C++-language backslash-escaping (meta character) conventions. For example, to include a backslash character in a delimiter string or in a character data-value, enter it as: \\.

Syntax:

**RadAttrFieldName** = *Radius-Attribute-Name*@"..."

#### **MVA-ByteString-Concat**

The MVA-ByteString-Concat extension treats the potentially multiple values of **Radius-Attribute-Name** as byte strings. It concatenates them using a length and value representation: every value is represented by a one-byte length header that identifies the length of the data value in the range 1–253 and is followed by the data value. The concatenation terminates with a one-byte 0-value trailer that indicates a data of length 0, such as the end of the concatenated byte-string. Ensure that values of the **Radius-Attribute-Name** attribute are of the natural BYTE datatype. If they are not, the appropriate unnatural conversion is applied.

**NOTE:** The **RadAttrFieldName** field must be declared in the CST schema to be of [VAR]BINARY type. If the **RadAttrFieldName** field is of a type other than [VAR]BINARY, Steel-Belted Radius Carrier reports an error.

Syntax:

**RadAttrFieldName** = *Radius-Attribute-Name*@\*

### Using MVA Facilities

The syntax of the MVA facilities supports single-valued attributes and multi-valued attributes. For single-valued attributes, @1, @^, @\$, and no-@-extension returns the value of the attribute. The @"..." and @\* extensions return one item in the concatenation; no delimiter appears for @"...", while the 0-byte terminator immediately follows the first item in the list for @\*.

Any field-to-attribute mapping in **sessionTable.ini** (whether its **Radius-Attribute-Name** is single-valued or multi-valued) that does not contain an MVA @ extension is considered to be equivalent to the same field-to-attribute mapping with an @^ extension appended to it.

In both the @"..." and the @\* cases, the individual values in the MVA are listed in the order they occur in the RADIUS packet. According to the RADIUS protocol, attribute sequence is significant.

Ensure that the size of a CST field is large enough to hold all content that may be generated, especially for MVA-concatenated fields (@"..." and @\*). One option you can use to solve the problem is to combine the MVA-concatenation facilities with the MVA-count facility (@#).

From the natural/unnatural design viewpoint, the MVA facilities lie somewhere in between; they can best be described as semi-natural, in two different ways:

- The @#, @N, @^, and @\$ facilities, which deal only with single values, can be considered to be both natural and unnatural:
  - They are natural in the sense of the previously defined natural RADIUS datatype SQL datatype mappings.
  - They are unnatural in the sense that they can be subjected to the unnatural RADIUS datatype SQL datatype mappings (the same way that no-@-extension can).
- The @"..." and @\* facilities, which deal with true multi-values, are necessarily non natural, because a concatenation operation does not exist in SQL. This leads to degrees of semi-naturalness:
  - In the case of CHAR values, it is (semi-)natural to separate concatenated items with a delimiter (field separator). However, the choice of the delimiter itself is conventional (unnatural).

- In the case of BYTE values, it is (semi-)natural to format the concatenated items in header/footer format, but there is no canonical way to do that because there are many protocol transfer syntaxes in existence – consider the difference between the format we support versus ASN1/BER, for example.
- In the case of INT/ADDR/TIME values, there is no well-defined concept of concatenation. MVAs in this case become unnatural to some extent, but it does seem natural to first convert the INT/ADDR/TIME values to CHAR/BYTE values, and then to concatenate those together.

## Admin Private Fields

You can add Private fields to the CST. Private fields are not processed by Steel-Belted Radius Carrier, and can include any information captured by or relevant to tools in the environment.

The Admin Private Fields section of **CurrentSessions.sql** looks like this:

```
#-----
ADMIN PRIVATE FIELDS
#-----
#PrivateField1          INT UNSIGNED
                        DEFAULT NULL,
#PrivateField2          VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci
                        DEFAULT NULL,
#PrivateField3          BINARY(32)
                        DEFAULT NULL,
```

## System Keys and Indexes

This section of the file contains keys and indexes that are required for system operation.



**CAUTION:** We strongly recommend that you not edit this section of the file. If any required key is not present in the CST, functions that depend upon it fail.

The System Keys and Indexes section of **CurrentSessions.sql** looks like this:

```
#-----
SYSTEM KEYS/INDEXES
```

```
#-----
PRIMARY KEY USING HASH          (Sbr_UniqueSessionId),
INDEX Sbr_NasName_Idx          (Sbr_NasName),
INDEX Sbr_ExpirationTime_Idx   (Sbr_ExpirationTime),
INDEX AcctMultiSessionId_Idx   (AcctMultiSessionId),
INDEX Sbr_MobileCorrelationId_Idx (Sbr_MobileCorrelationId),
INDEX Sbr_AcctSessionId_Idx   (Sbr_AcctSessionId),
INDEX Sbr_UserName_Idx        (Sbr_UserName),
INDEX Sbr_CallingStationId_Idx (Sbr_CallingStationId),
INDEX Sbr_CalledStationId_Idx (Sbr_CalledStationId),
INDEX Sbr_Ipv4Address_Idx      (Sbr_Ipv4Address),
INDEX Sbr_TransactionId_Idx    (Sbr_TransactionId )
```

## Admin Keys and Indexes

This section of the file contains keys and indexes that are required for system operation.



**CAUTION:** We strongly recommend that you not edit this section of the file; it contains system-generated entries. If any required key is not present in the CST, functions that depend upon it fail.

The System Keys and Indexes section of **CurrentSessions.sql** looks like this:

```
#-----
ADMIN KEYS/INDEXES
#-----
#,
INDEX Some_Idx USING HASH (What, Ever)
)
ENGINE = ndbcluster NOTE: CreatedB.sh fiddles with this line!
;
```

```
#=====
```

## SSR Datatypes

This section describes the datatypes that Session State Register supports.

**NOTE:** We do not support any datatype not explicitly addressed in this document. This applies to datatypes themselves, sizes of datatypes, character sets, and their collations.

### Definitions

Data in a RADIUS attribute may be one of six datatypes:

- **NULL:** No data at all; absence of data.
- **INT:** Data that is either a bare (numerical, counting) measure or a semantic (enumeration, ordinal) value, representing one value out of a specific set of values.

Examples: 42 (measuring a person's age in years); 0 (as a stand-in for false).

- **ADDR:** IPv4 addresses. These are ultimately 32-bits quantities as far as networks are concerned, but are usually represented as integers or as dotted-quad strings. Example:  
10101100000110010110000101001101 = 0xac19614d = 2,887,344,461 = 172.25.97.77.
- **TIME:** Seconds elapsed in the UNIX epoch, commencing at 00:00:00 (the stroke of midnight) on the morning of January 1, 1970 UTC. Example: 1,142,370,727 = 2006-03-14T16:12:07-0500 = 4:12:07PM Eastern Daylight Time, March 14, 2006.
- **CHAR:** Information expressed in natural languages. Example: *Administration Building*.
- **BYTE:** Encoded data. Example: Encrypted communications.

Table 25 on page 172 summarizes these data types:

Table 25: Terminology for Datatypes

This Document	RADIUS Spec	SQL	Core SBR & SBR Dictionaries	SBR/SSR Implementation
NULL	(absent)	NULL	RAD_TYPE_CONSTANT	(absent)

Table 25: Terminology for Datatypes (continued)

This Document	RADIUS Spec	SQL	Core SBR & SBR Dictionaries	SBR/SSR Implementation
INT	integer (32-bit unsigned, big-endian)	TINYINT SMALLINT MEDIUMINT INT  SIGNED or UNSIGNED	RAD_TYPE_INTEGER, RAD_TYPE_INT1 RAD_TYPE_HEX4 RAD_TYPE_IPADDR_POOL integer signed-integer int1 int4 hex1 hex4 ipxaddr-pool	FieldFormat_UInteger, FieldFormat_SInteger
ADDR	address (32-bit)	INT  UNSIGNED	RAD_TYPE_IPADDR, RAD_TYPE_IPADDR_POOL ipaddr ipaddr-pool	FieldFormat_UInteger
TIME	time (32-bit)	TIMESTAMP	RAD_TYPE_TIME time	FieldFormat_UInteger
CHAR	text (utf8, 0-255 bytes)	CHAR VARCHAR	RAD_TYPE_STRING RAD_TYPE_STRINGNZ string stringnz	FieldFormat_String, FieldFormat_StringFixed
BYTE	string (0-255 bytes)	BINARY VARBINARY	RAD_TYPE_HEXSTRING, RAD_TYPE_IPV6ADDR, RAD_TYPE_IPV6PREFIX, RAD_TYPE_IPV6INTERFACE hexadecimal ipv6addr ipv6prefix, ipv6interface	FieldFormat_Raw, FieldFormat_RawVar

### System Field Datatypes

The datatypes for the System fields have already been specified as listed in [Table 25 on page 172](#). Some are fixed, and some are customizable.

# RadAttr Fields

A RadAttr field declared in `sessionTable.ini` as `RadAttrFieldName = Radius-Attribute-Name` associates a CST field (`RadAttrFieldName`) with a RADIUS attribute (`Radius-Attribute-Name`), that involves two datatypes:

- The RADIUS attribute (`Radius-Attribute-Name`) has a RADIUS datatype specified by the dictionary in which Steel-Belted Radius Carrier finds it.
- The CST field (`RadAttrFieldName`) has a SQL datatype determined by the field's declaration in the CST schema (`CurrentSessions.sql`).

These two datatypes must match or be compatible with one another so the RADIUS datatype can be mapped or converted to the SQL datatype. In this context:

- **Match**—A natural /primary/canonical/favorite mapping, as specified in [Table 26 on page 174](#).
- **Compatible**—An unnatural/secondary/conventional/formatting conversion, as specified in [Table 26 on page 174](#).

Take care to coordinate datatypes to avoid processing errors.

## Matching/Natural Datatype Mappings

[Table 26 on page 174](#) lists the natural datatype conversions from RADIUS datatypes to SQL datatypes, together with the usual way to display values stored in the database using the SQL SELECT command as implemented in `ShowSessions.sh`.

Table 26: Matching/Natural Mappings

RADIUS Datatype	SQL Datatype	SQL SELECT Display Command
(absent)	NULL	Conventionally displayed in Steel-Belted Radius Carrier as <code>(n u l l)</code> (9 characters, 3 of which are spaces,).
integer	INT UNSIGNED	Supported natively, for example (with auxiliary formatting options, such as commas for thousands separators, available):  <code>SELECT RadAttrFieldName FROM CurrentSessions;</code>  In the case of enumerative values, you should use the <code>CONCAT()</code> built-in function to display the name of the enumerated value, followed by the integer value in parentheses.



Table 26: Matching/Natural Mappings (continued)

RADIUS Datatype	SQL Datatype	SQL SELECT Display Command
address	INT UNSIGNED	<p>The standard dotted-quad string notation for IPv4 addresses is available using the <b>INET_NTOA()</b> built-in function:</p> <p><b>SELECT INET_NTOA (<i>RadAttrFieldName</i>) AS <i>RadAttrFieldName</i> FROM CurrentSessions;</b></p>
time	TIMESTAMP	<p>Supported natively, for example (with auxiliary formatting options available):</p> <p><b>SELECT <i>RadAttrFieldName</i> FROM CurrentSessions;</b></p>
text	[VAR]CHAR	<p>Supported natively. Make the CHAR nature of the data explicit by surrounding it with double quotation marks:</p> <p><b>SELECT CONCAT("\", <i>RadAttrFieldName</i> , "\") AS <i>RadAttrFieldName</i> FROM CurrentSessions;</b></p>
string	[VAR]BINARY	<p>Supported natively. Make the BYTE nature of the data explicit by surrounding it with single quotation marks:</p> <p><b>SELECT CONCAT("\", LCASE(HEX(<i>RadAttrFieldName</i> )) , "\") AS <i>RadAttrFieldName</i> FROM CurrentSessions;</b></p>

## Sample ShowSessions.sh Report

Figure 14 on page 175 shows sample output of **ShowSessions.sh**:

```
CurrentSessions:
+ -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- + (1)
CORE
    UniqueSessionId: 'e86cb6d3717ab5331c68e6bf2b8ec6ef'x
      CreationTime: 2017-08-17 12:44:58 (TZ=+00:00)
      ExpirationTime: 2017-08-18 12:44:58 (TZ=+00:00)
      Ipv4Address: 10.212.10.17
      IpAddrPool: (n u l l)
      NasName: "LOCALHOST"
```



- The Core, Feature, Optional, RadAttr, and Private fields are grouped into separate titled sections, for the convenience of the reader. You can change this or remove titles by modifying the **ShowSessions.sh** script.
- The System field names have been made easier to read by removing their Sbr\_ prefixes, and the displayed names of the Admin fields have been customized to include the prefix **My**.
- The various kinds of SQL SELECT displays mentioned in [Table 26 on page 174](#) are shown:
  - NULL values are shown for 10 fields. To eliminate fields that are always NULL (such as disabled Feature fields), comment them out in **ShowSessions.sh** so they do not appear in the printout.
  - Raw INT data for the NasPort and MyRadAttrField1 fields is displayed. (Note the formatting of the separators in NasPort. You can customize this.)
  - Enumerative INT data for the IpAddrPool, Status, and NasPortType fields is displayed.
  - ADDR data is displayed for the Ipv4Address field.
  - TIME data is displayed for the CreationTime and ExpirationTime fields.
  - CHAR data is displayed for the NasName, UserName, AcctSessionId, and MyRadAttrField2 fields.
  - BYTE data is displayed for the UniqueSessionId, TransactionId, and MyRadAttrField3 fields.

Concerning the definitions of the RadAttr fields in this example:

- The SQL declaration of the RadAttr fields is:

```
RadAttrField1 INT UNSIGNED DEFAULT NULL,
RadAttrField2 VARCHAR(24) CHARSET utf8 COLLATE utf8_general_ci DEFAULT NULL,
RadAttrField3 BINARY(24) DEFAULT NULL,
```

- The **sessionTable.ini** file is:

```
[AuthRequest]
RadAttrField1 = NAS-Port
RadAttrField2 = User-Name
RadAttrField3 = User-Password
[AuthResponse]
[AcctRequest]
[AcctResponse]
```

The definition of the RadAttrField1 has the same value as the NAS-Port field, though the latter has some formatting applied to it (at the natural display level, not at the unnatural storage level; see the following subsection).

The definition of RadAttrField2 is the same as User-Name, except that the latter has been processed by Steel-Belted Radius Carrier (as mentioned previously).

The definition of RadAttrField3 demonstrates the kind of padding that occurs with SQL BINARY fields. Had it been declared VARBINARY(24), the padding zeros would not have appeared.

**NOTE:** The RADIUS User-Password attribute is encrypted (hence, BYTE) data.

### Compatible/Unnatural Datatype Conversions

In addition to the matching/natural mappings defined previously, Steel-Belted Radius Carrier SSR supports the compatible/unnatural conversions listed in [Table 27 on page 178](#).

**Table 27: Compatible/Unnatural Conversions**

RADIUS Datatype	SQL Datatype			
	INT UNSIGNED	TIMESTAMP	[VAR]CHAR	[VAR]BINARY
<b>integer</b>	Natural Match	Coerce	Base10	Base256
<b>address</b>	Natural Match	Coerce	Dotted Quad	Base256
<b>time</b>	Coerce	Natural Match	ISO8601	Base256
<b>text</b>	Length	None	Natural Match	Coerce
<b>string</b>	Length	None	Hex	Natural Match

This list further explains the values in [Table 27 on page 178](#):

- **Coerce**—Re-interpret the underlying bits of the RADIUS datatype as if they were the underlying bits of the corresponding SQL datatype.
- **Length**—Length in bytes (not in characters in the CHAR case). A terminating NULL byte, if any, is not counted. This conversion is supported for its helpfulness in checking database-storage padding and truncation issues.
- **Base10**—Express an integer in the usual decimal notation, using the Unicode Basic Latin characters 0...9, possibly preceded by a minus-sign.
- **Dotted Quad**—Conventional notation for IPv4 addresses.
- **ISO8601 (RFC 3339)**—Conventional international standard format YYYY-MM-DDThh:mm:ssZ. Both T and Z are required elements. The Z indicates UTC; if it is not present, the interpretation defaults to local time zone.

- Hex—Conventional mapping of BYTE strings to hex-strings (character-strings of even length, using only the characters [0-9a-fA-F]).
- Base256—Conventional mapping of 32-bit (big-endian) INT values to 4-byte BYTE strings: write the given number in the form  $N = A \cdot 256^3 + B \cdot 256^2 + C \cdot 256 + D$ , then use the sequence A/B/C/D as the base-256 encoding, suppressing leading zeros. This is, in essence, a kind of coercion, if we posit a 4-byte big-endian environment.
- None—The potential conversion is not supported.

**NOTE:** There is a major conceptual difference between the preceding section and this one. We highlight the address>dotted-quad conversion for the sake of concreteness, but the argument applies equally well to integer>base-10, time>ISO8601, and string>Hex.

In the natural mechanism, IPv4 addresses are stored in the CST as INT UNSIGNED values, and the dotted-quads are merely used as a device for displaying the values to users. This is how the Sbr\_Ipv4Address System field is implemented. But in the unnatural case, the IPv4 address is stored in the CST itself as dotted-quad CHAR data.

Other things being equal, the natural scheme is preferred because SQL's extensive built-in facilities and stored procedures are adequate for all data-manipulation and data-display tasks, and are much more flexible than our unnatural scheme. The unnatural scheme is supported with a certain class of tools that foster datatype conflation. This is traditional for SQL-based programming, and SQL itself includes convenience tools of this class. For example, MySQL permits the (64-bit) number 20060314161207, the character-string 20060314161207, and the timestamp 2006-03-14 16:12:07 to be used interchangeably.

## Customizing Datatypes and Data Sizes

You have some latitude for customizing datatypes (including character sets and collations) and sizes in many System and RadAttr CST fields. This section describes the rules for modification.

To understand the problem, consider the **Sbr\_NasName** field, but pretend for simplicity that it always contains the contents of a NAS-Identifier RADIUS attribute. According to RFC 2865, the NAS-Identifier attribute has a BYTE datatype, and can have any size from 0 through 255 bytes. Consequently, you can safely declare the **Sbr\_NasName** field to be [VAR]BINARY(255). However, in most environments, NAS-Identifiers are relatively short and human-readable; that data value is the underlying byte-string of an encoded character-string. In such environments, declaring **Sbr\_NasName** to be [VAR]BINARY(255) is wasteful in terms of database storage and network bandwidth used and less useful because the resulting string is unreadable. In such cases, you might declare **Sbr\_NasName** to be VARCHAR(24) using UTF8 text strings.

- For System Core and Feature fields, use a datatype and size that guarantees full fidelity of information (no semantic data loss). If you use an inappropriate datatype or size, Steel-Belted Radius Carrier SSR can fail in unpredictable ways because the data in question is critical to stable processing. If the data in a System Core or Feature field is unreliable, Steel-Belted Radius Carrier might not work correctly.
- If Steel-Belted Radius Carrier does not depend on the data in question, a less restrictive rule applies. For example, consider the **Sbr\_UserName** field. Again, the RADIUS RFC guarantees that [VAR]BINARY(255) is safe. You can declare the **Sbr\_UserName** field as VARCHAR(24) using Latin1 text strings, but might lead to problems if users want to use non-Latin characters (for example, Japanese or Devanagari) in user names.

To apply these rules, be familiar with the RADIUS RFCs and the network environment.

**NOTE:** If you have included a customizable field for the Delegated-IPv6-Prefix attribute in MySQL table, its data type needs to be set as VARBINARY(16) for displaying the field in IPv6 address format in the **SessionControl.sh** script output.

## Character Sets and Collations

The SSR database supports many character sets (*charsets*). Steel-Belted Radius Carrier SSR sites can use any charset and any associated collation supported by the SSR database that does not use the NULL byte in its encoding. Examples of permissible charsets are: ASCII, Latin1, and UTF8. A non-example is UCS2 (which is the straight-forward 2-byte encoding of the Basic Multilingual Plane (BMP), Unicode U+0000...U+FFFF).

Because charsets in the SSR database can have an expansion factor (1, 2, or 3) associated with them, specifications for field sizes can be misleading. For example, if a field in the CST schema is declared **VARCHAR(7) CHARSET utf8**, then the field is guaranteed to hold up to seven Unicode characters encoded in UTF8. Because UTF8 has an expansion factor of 3, a VARCHAR(7) field can hold up to  $3 \times 7 = 21$  bytes of information.

## Truncation and Padding

The SQL datatypes **[VAR]CHAR(N)** and **[VAR]BINARY(N)** reserve a fixed number of bytes, called the field's storage size, for every stored value in NDB. The storage size for **[VAR]CHAR** is the charset expansion factor times N; the storage size for **[VAR]BINARY** is N. Steel-Belted Radius Carrier always presents exactly the full storage size of data to NDB.

In cases where a field value is less than the storage size, Steel-Belted Radius Carrier pads the value out to the full storage size before presenting it to NDB for storage. The padding value is the SPACE character in the **[VAR]CHAR** case and the NUL byte in the **[VAR]BINARY** case.

In cases where a field value is greater than the storage size, Steel-Belted Radius Carrier truncates it to the field storage size. No log file messages are written if data is truncated.

The maximum storage size for any field in the CST is 4096 bytes, large enough for any System field or natural RadAttr field. This is sufficient to store all the raw data carried in any RADIUS attribute value, even if the whole RADIUS packet consists of one large multi-value attribute.

## Small Integers

Truncation and padding are not used with INT numerical data. Because the RADIUS specification discusses only 4-byte INT data, you can always safely declare a **RadAttrFieldName** in the CST schema as an INT.

If you are certain that, in your network, all values of **RadAttrFieldName** fit into a TINYINT (1 byte), SMALLINT (2 bytes), or MEDIUMINT (3 bytes), you can declare **RadAttrFieldName** to be of an appropriate type in the SQL schema. However, if an illegal overflow value is presented to Steel-Belted Radius Carrier for storage in a field that is too small, the value that Steel-Belted Radius Carrier stores is the largest value that the field will hold (0xff = 255, 0xffff = 65,535, 0xffffffff = 16,777,215 respectively). No log file message is generated. As a general rule, do not use TINYINT, SMALLINT, or MEDIUMINT unless you are sure the network can tolerate this potential loss of information.

## Signed Integers

In the original RADIUS specifications, all integer attributes were specifically unsigned. Recent usages have started to use signed attributes. For this reason, our implementation allows **RadAttrFieldNames** to be specified as SIGNED fields in the CST schema. Steel-Belted Radius Carrier SSR supports this only as a natural facility: you must declare the sign of SQL INT datatypes correctly. There are no (natural or unnatural) conversions between signed and unsigned integers.

For signed INT quantities, you can use TINYINT, SMALLINT, or MEDIUMINT in SSR database schema declarations if you know the values in your network can fit into these smaller-sized fields. If a value too large for a field is presented to SBR Carrier for storage, SBR Carrier stores the largest possible value (that is, most-positive, which is 0x7f = 127, 0x7fff = 32,767 or 0x7fffffff = 8,388,607) in the case of overflow, and the smallest possible value (that is, the most-negative, which is -128, -32,768 or -8,388,608 respectively) in the case of underflow. No log file message is generated.

## Unsupported SQL Datatypes

RadAttr CST fields do not support SQL datatypes other than those discussed in this section (INTs of various sizes and sign, CHAR, VARCHAR, BINARY, and VARBINARY).

Unsupported datatypes include INT datatypes such as BIGINT, BIT, DECIMAL, TIME and DATE; CHAR datatypes such as TEXT; and BYTE datatypes such as BLOB.

# Other Session State Register Configuration Files

## IN THIS CHAPTER

- Tuning Other SSR Tables Overview | 182
- Configuring the `dbclusterndb.gen` File | 182
- Customizing Other Tables and Scripts | 193
- Managing IP Pools | 194

This chapter presents information about maintaining and working the Session State Register module (SSR). This chapter covers the following topics:

## Tuning Other SSR Tables Overview

Several other files control tables in the SSR database. These usually do not require modification because most values are system-generated, but some settings may be changed. In particular, the Management Mode field in the `dbclusterndb.gen` file is often changed to put SBR Carrier nodes into management mode for maintenance. (See the section on *Session State Register Administration* in the *SBR Carrier Administration and Configuration Guide*.) Most other settings in all files are generated by the system and do not usually require changes.

## Configuring the `dbclusterndb.gen` File

Use the `dbclusterndb.gen` file to configure the database settings used by each SBR Carrier node to access the SSR database. The `dbclusterndb.gen` file is stored in the `radiusdir` directory, usually `/opt/JNPRsbr/radius`.

### [Bootstrap] Section

The [Bootstrap] section ([Table 28 on page 183](#)) of `dbclusterndb.gen` specifies information that SBR Carrier nodes use to load SSR functions.



```
[Bootstrap]
LibraryName = dbclusterndb
Enable = 1
ManagementMode = 0
```

**Table 28: dbclusterndb.gen [Bootstrap] Fields**

Parameter	Function
LibraryName	<p>Specifies the name of the cluster database module. Default value is <b>dbclusterndb</b>.</p> <p>Do not change this unless you are advised to do so by Juniper Networks Technical Support.</p>
Enable	<ul style="list-style-type: none"> <li>• If set to 0, the high availability functionality is disabled.</li> <li>• If set to 1, the high availability functionality is enabled.</li> </ul> <p>Default value is 1 in the file provided with SSR. If this setting is removed from the <b>dbclusterndb.gen</b> file, default value switches to 0.</p>
ManagementMode  For information about management mode, see the section on <i>Session State Register Administration</i> in the <i>SBR Carrier Administration and Configuration Guide</i> .	<ul style="list-style-type: none"> <li>• If set to 0, the SBR Carrier server operates in standard high-availability mode.</li> <li>• If set to 1, the SBR Carrier server operates in management mode.</li> </ul> <p>Default value is 0.</p>

## [NDB] Section

The [NDB] section ([Table 29 on page 184](#)) of **dbclusterndb.gen** identifies how SBR Carrier nodes access the SSR database.

```
[NDB]
ManagementNode = 127.0.0.1:5235;nodeid=30
ConnectRetries = 3
DelayBetweenConnectRetriesSec = 5
TimeoutForFirstAliveSec = 10
WaitForAllNodesAlive = 0
TimeoutAfterFirstAliveSec = 10
NdbHandles = 32
NdbHandlesAlert = 1
```

NDBHardErrorThreshold = 10

Table 29: dbclusterndb.gen [NDB] Fields

Parameter	Function
ManagementNode	<p>Specifies the NDB connect-string value, made up of the IP address of the management node hosts, the port the management node uses for connection requests, and the node ID (<b>NDB connect-string</b>) of the local SBR Carrier node.</p> <p>This information is created by the installation script and should only be changed by that script or under direction of JTAC.</p>
ConnectRetries	<p>Specifies how many times SBR Carrier tries to connect to the management nodes.</p> <p>Default value is 3.</p>
DelayBetweenConnectRetriesSec	<p>Specifies how many seconds SBR Carrier waits between retries when trying to connect to the management nodes.</p> <p>Default value is 5 seconds.</p>
TimeoutForFirstAliveSec	<p>Specifies how many seconds SBR Carrier waits for the first NDBD server to confirm that it can communicate before communicating with the database cluster.</p> <p>Default value is 10 seconds.</p>
WaitForAllNodesAlive	<ul style="list-style-type: none"> <li>• If set to 0, SBR Carrier does not wait for confirmation that all NDB nodes are alive before communicating with the database cluster.</li> <li>• If set to 1, SBR Carrier waits for confirmation that all NDB nodes are alive before communicating with the database cluster.</li> </ul> <p>Default value is 0.</p>
TimeoutAfterFirstAliveSec	<p>Specifies how many seconds SBR Carrier waits after the first NDBD server alive indicator.</p> <p>Default value is 10 seconds.</p>

Table 29: dbclusterndb.gen [NDB] Fields *(continued)*

Parameter	Function
NdbHandles	<p>Specifies the number of NDB handles used for parallel database transactions. One NDB handle is needed for each ongoing database operation. Therefore, the number of available handles poses an upper limit to the number of concurrent operations a SBR Carrier server can make to NDB. Performance increases with concurrency up to a point, where thread overhead overwhelms the benefits of concurrency. The best balance depends on your environment.</p> <p><b>NOTE:</b> Each NDB handle uses more than 32K of memory. Increasing the value of NdbHandles increases the start time for SBR Carrier very slightly.</p> <p>Enter a value in the range 1–128. Note that one handle is attached permanently to each cache thread, and one handle is reserved by the system for special usage when in ManagementMode.</p> <p><b>NOTE:</b> The NdbHandles setting, and its associated alert message, only counts operational handles, i.e., not those reserved for MgmtMode or for caching threads.</p> <p>Default value is 32.</p>
NdbHandlesAlert	<ul style="list-style-type: none"> <li>• If set to 0, the SBR Carrier server does not record NDB handle messages in its log file.</li> <li>• If set to 1, the SBR Carrier server records messages identifying the maximum number of concurrent NDB handles that have been used since the server was restarted in its log file. The messages take the form: <p style="text-align: center;"><b>Max concurrent NDB handles = N</b></p> </li> </ul> <p>Default value is 0.</p>

Table 29: dbclusterndb.gen [NDB] Fields (continued)

Parameter	Function
NDBHardErrorThreshold	<p>Specifies the threshold value for NDB hard errors. If the number of hard errors exceeds the threshold value, SBR Carrier starts to monitor the cluster health and determines whether to persist sessions in the local file on the SBR Carrier server on the basis of the value set in the <b>FallbackLocal</b> parameter in the <b>radius.ini</b> file.</p> <p>You can enter the value in the range from 1 through 100. Default value is 10.</p> <p>This parameter is reloaded every time when SBR Carrier receives a hang-up (HUP) signal.</p>

**[Database] Section**

The [Database] section ([Table 30 on page 186](#)) of **dbclusterndb.gen** controls how SBR Carrier front ends accesses the SSR database.

```
[Database]
Database = SteelBeltedRadius
Retries = 6
DelayBetweenRetriesMillisec = 50
RetryAlertThreshold = 4294967295
```

Table 30: dbclusterndb.gen [Database] Fields

Parameter	Function
Database	<p>Specifies the name of the database used by SBR Carrier.</p> <p>Default value is <b>SteelBeltedRadius</b>.</p>
ReconnectOnHUP	<p>Controls whether the database cluster is disconnected and reconnected after a HUP signal is sent to the dbcluster plugin.</p> <ul style="list-style-type: none"> <li>• If set to 1, the database reconnects after receiving a HUP signal.</li> <li>• If set to 0, the database does not reconnect after receiving a HUP signal.</li> </ul> <p>Default value is 0.</p>

Table 30: dbclusterndb.gen [Database] Fields (continued)

Parameter	Function
Retries	<p>Specifies how many times SBR Carrier tries to connect to the management nodes before giving up.</p> <p>Default value is 6.</p>
DelayBetweenRetriesMillisec	<p>Specifies the base number of milliseconds SBR Carrier waits before retrying a database operation. The first retry delay is 1 x the value specified for DelayBetweenRetries, the second retry delay is 2 x the value specified for DelayBetweenRetries, and the nth retry delay is n x the value specified for DelayBetweenRetries.</p> <p>Default value is 50 milliseconds.</p>
RetryAlertThreshold	<p>Specifies the threshold for recording log messages when a retry that exceeds the threshold is attempted. The log message identifies why the retry was attempted (that is, why the previous attempt failed).</p> <ul style="list-style-type: none"> <li>• If RetryAlertThreshold is set to 0, log messages are written before every database retry.</li> <li>• If RetryAlertThreshold is set to the value of (Retries -1), log messages are written before the last retry and after the last retry (if it fails).</li> <li>• If RetryAlertThreshold is set to the value of Retries, log messages are written only after the last retry fails (that is, when the last retry failed and no further attempt will be made).</li> <li>• If RetryAlertThreshold is set to a value of greater than the value of Retries, log messages are not recorded.</li> </ul> <p>Default value is 4294967295 (0xffffffff).</p>
UseConnectionManager	<p>UseConnectionManager=&lt;bool&gt;. This parameter should always be set to true unless a change is recommended by Juniper Networks Technical Support. This might have a minor performance impact on throughput, bandwidth-bound installations. Contact your sales engineer or Juniper Networks Technical Support for more information.</p> <p>The default is enabled.</p>

### [IpAddressPools] Section

The [IpAddressPools] section ([Table 31 on page 188](#)) of **dbclusterndb.gen** specifies aging and caching parameters for IP address pools. Some of these settings are system-wide and cannot be overridden. Other settings establish system defaults, which can be overridden for specific IP address pools.

```
[IpAddressPools]
MinUnusedAgeSec = 300
MaxAgeRetries = 2
AgePercent = 50
AgeRetryAlertThreshold = 4294967295
NumCacheThreads = 2
StartCachingAtBootTime = 1
CacheThreadSleepMin = 1000
CacheThreadSleepMax = 2000
CacheLowWater = 100
CacheHighWater = 250
CacheChunkSize = 50
EmergencyChunkSize = 1
CacheAlertThreshold = 0
CacheThreadVerbose = 0
```

**Table 31: dbclusterndb.gen [IpAddressPools] Fields**

Parameter	Function
MinUnusedAgeSec	<p>Specifies how many seconds an IP address can remain unused before SBR Carrier reassigns it.</p> <p>Default value is 300 seconds.</p>
MaxAgeRetries	<p>Specifies the number of times SBR Carrier attempts to retrieve acceptably aged IP addresses before retrieving any available address.</p> <p>If MaxAgeRetries is 0, SBR Carrier makes one attempt to look for addresses that have been idle for at least MinUnusedAgeSec seconds, then no retries if enough addresses are not found.</p> <p>Default value is 2.</p>

Table 31: dbclusterndb.gen [IPAddressPools] Fields (continued)

Parameter	Function
AgePercent	<p>Specifies the percentage (a number in the range 0–100) that SBR Carrier uses as the multiplier for MinUnusedAgeSec when enough IP addresses that have been idle for MinUnusedAgeSec cannot be found. For example, if MinUnusedAgeSec is 400 and AgePercent is 75, SBR Carrier would look for addresses idle for at least 400 seconds, then look for addresses that have been idle for at least 300 (<math>400 * 0.75</math>) seconds, then look for addresses that have been idle for at least 225 (<math>300 * 0.75</math>) seconds.</p> <p>If AgePercent is 100, only the original MinUnusedAgeSec is used.</p> <p>If AgePercent is 0, then age is disregarded.</p> <p>Default value is 50.</p>
AgeRetryAlertThreshold	<p>Specifies the threshold (0–4294967295) for recording log messages when an aged-based retry that exceeds the threshold is attempted. The log message identifies why the retry was attempted (that is, why the previous attempt failed).</p> <ul style="list-style-type: none"> <li>• If AgeRetryAlertThreshold is set to 0, log messages are written before every age-based retry.</li> <li>• If AgeRetryAlertThreshold is set to the value of (MaxAgeRetries - 1), log messages are written before the last age-based retry and after the last retry (if it fails).</li> <li>• If AgeRetryAlertThreshold is set to the value of MaxAgeRetries, log messages are written only after the last retry fails (that is, when no addresses have been found and no further attempt will be made).</li> <li>• If AgeRetryAlertThreshold is set to a value of greater than the value of MaxAgeRetries, log messages are not recorded.</li> </ul> <p>Default value is 4294967295 (0xffffffff).</p>
NumCacheThreads	<p>Specifies how many parallel threads SBR Carrier uses to cache IP addresses.</p> <p>Default value is 2.</p>

Table 31: dbclusterndb.gen [IPAddressPools] Fields (continued)

Parameter	Function
StartCachingAtBootTime	<ul style="list-style-type: none"> <li>• If set to 1, SBR Carrier fills its IP address pool cache immediately when it is rebooted.</li> <li>• If set to 0, SBR Carrier fills its IP address pool cache when it receives an address request.</li> </ul> <p>Default value is 1.</p>
CacheThreadSleepMin	<p>Specifies the minimum range of time (0–4294967295) a cache-filling thread waits before it goes to the database to get another cache of IP addresses. Set this parameter to better manage your caching threads.</p> <p>Default value is 1000 milliseconds.</p>
CacheThreadSleepMax	<p>Specifies the maximum range of time (<i>CacheThreadSleepMin</i>–4294967295) a cache-filling thread waits before it goes to the database to get another cache of IP addresses. Set this parameter to better manage your caching threads.</p> <p>Default value is <math>2 * \text{CacheThreadSleepMin}</math> milliseconds.</p>
CacheLowWater	<p>Specifies the minimum number of addresses that must be available in the address cache for an IP address pool. When the number of addresses in a server's cache falls below the <i>CacheLowWater</i> value, the server begins requesting blocks of IP addresses</p> <p>Default value is 100.</p>
CacheHighWater	<p>Specifies the number of addresses that must be available in a server's IP address cache for an IP address pool before it stops adding addresses to the cache</p> <p>The <i>CacheHighWater</i> value must be greater than or equal to the <i>CacheLowWater</i> value.</p> <p>Default value is 250.</p>
CacheChunkSize	<p>Specifies the number of addresses to retrieve every time SBR Carrier requests a block of IP addresses for an IP address pool.</p> <p>Default value is 50.</p>



Table 31: dbclusterndb.gen [IPAddressPools] Fields (continued)

Parameter	Function
EmergencyChunkSize	<p>Specifies the (0–CacheChunkSize) number of addresses to retrieve every time <i>SBR Carrier</i> requests a block of IP addresses to use from a pool, when that pool's cache is empty, to directly retrieve those addresses from the IP address table in the database and put them in the cache.</p> <p>Default value is 1.</p>
CacheAlertThreshold	<p>Specifies the threshold (0–4294967295) for recording log messages when the number of addresses in the cache falls below the threshold value.</p> <ul style="list-style-type: none"> <li>• If CacheAlertThreshold is set to 0, log messages are not written.</li> <li>• If CacheAlertThreshold is set to 1, log messages are written when the address cache is empty.</li> <li>• If CacheAlertThreshold is set to a value equal to or greater than the sum of the values of CacheHighWater and CacheChunkSize, log messages are written whenever an address is pulled from the cache.</li> </ul> <p>Default value is 0.</p>
CacheThreadVerbose	<ul style="list-style-type: none"> <li>• If set to 1, print out an informational message about what the caching thread just did. The message contains: thread identity (OS thread ID number), how long it napped, pool cached, number of IP addresses cached, and the length of time it took to retrieve the addresses from the database.</li> <li>• If set to 0, no caching thread information is recorded.</li> </ul> <p>Use this parameter to fine-tune the caching parameters.</p> <p>Default value is 0.</p>

### [IPAddressPools:PoolName] Section

The [IPAddressPools:PoolName] section ([Table 32 on page 192](#)) of **dbclusterndb.gen** identifies how SBR Carrier specifies the override settings for IP address aging for a named IP address pool. You can create as many [IPAddressPools:PoolName] section as you require in the **dbclusterndb.gen** file to tune caching for individual address pools.

[IPAddressPools:Gold]

```

StartCachingAtBootTime = 0
CacheLowWater = 10
CacheHighWater = 25
CacheChunkSize = 5
EmergencyChunkSize = 2
CacheAlertThreshold = 1

```

**NOTE:** You cannot use an `[IpAddressPools:PoolName]` section to override settings other than those listed here. For example, you cannot enter a `MaxAgeRetries` setting in an `[IpAddressPools:PoolName]` section to override the value specified in the `[IpAddressPools]` section.

**NOTE:** Authentication requests may be rejected when you have configured the default settings for parameters in the `[IpAddressPools:PoolName]` section and the named IP address pool consists of a smaller number of IP addresses. To avoid this, we recommend that you set the number of IP addresses in the pool to a value greater than the values entered for **CacheLowWater**, **CacheHighWater**, and **CacheChunkSize**.

**Table 32: dbclusterndb.gen [IpAddressPools:PoolName] Fields**

Parameter	Function
StartCachingAtBootTime	<ul style="list-style-type: none"> <li>• If set to 1, SBR Carrier fills its IP address pool cache immediately when it is rebooted.</li> <li>• If set to 0, SBR Carrier fills its IP address pool cache when it receives an address request.</li> </ul> <p>Default value is 0.</p>
CacheLowWater	Specifies the minimum number of addresses that must be available in the named pool's IP address cache. When the number of addresses in the pool's cache falls below the CacheLowWater value, the server begins requesting blocks of IP addresses.
CacheHighWater	<p>Specifies the number of addresses that must be available in a server's IP address cache before it adding addresses to the IP address cache for the named IP address pool.</p> <p>The CacheHighWater value must be greater than or equal to the CacheLowWater value.</p>

Table 32: dbclusterndb.gen [IpAddressPools:PoolName] Fields (continued)

Parameter	Function
CacheChunkSize	Specifies the number of addresses to retrieve every time SBR Carrier requests a block of IP addresses for the named IP address pool.
EmergencyChunkSize	Specifies the (0– <i>CacheChunkSize</i> ) number of addresses to retrieve every time SBR Carrier requests a block of IP addresses to use from a pool, when that pool's cache is empty, to directly retrieve those addresses from the IP address table in the database and put them in the cache.  Default value is 1.
CacheAlertThreshold	Specifies the threshold (0–4294967295) for recording log messages when the number of addresses in the cache falls below the threshold value.  <ul style="list-style-type: none"> <li>• If CacheAlertThreshold is set to 0, log messages are not written.</li> <li>• If CacheAlertThreshold is set to 1, log messages are written when the address cache is empty.</li> <li>• If CacheAlertThreshold is set to a value equal to or greater than the sum of the values of CacheHighWater and CacheChunkSize, log messages are written whenever an address is pulled from the cache.</li> </ul> Default value is 0.

## Customizing Other Tables and Scripts

In addition to the Current Sessions Table (CST), SSR uses other SSR database tables. These tables may not be customized, except for the entries noted in [Table 33 on page 194](#):

Table 33: Other Customizable Database Fields

Table/Feature/Script	Customizable Field	Description	SQL Datatype & Size
<b>Sbr_IpPools</b> (in <b>IpAddr.sql</b> )	<b>Name</b> field	The <b>Name</b> field holds the names of the IPv4 Address Pools created by <b>AddPool.sh</b> script.	Defaults to VARCHAR(24) utf8  Customizable character set and size. Maximum size is 65535 chars.
<b>Sbr_UserConcurrency</b> (in <b>UserConcurrency.sql</b> )	<b>Id</b> field	The <b>Id</b> field holds an identifier used for SSR user concurrency feature.	Defaults to VVARCHAR(32) utf8  Customizable character set and size. Maximum size is 65535 chars.  The data type and size of this field must be the same as that of the <b>Sbr_UserConcurrencyId</b> field in the CST.
<b>Monitor.sh</b>	Date-stamp	An admin script that prints out a date-stamp that you can customize.	

## Managing IP Pools

The **AddPoolRangeBatch.pl** script can be used to manage IP pools in SBR Carrier. You can add IP Pools to the SSR database by running this script as the hadm user. The **AddPoolRangeBatch.pl** script is available in the **/opt/JNPRhadm/** directory.

Execute:

```
hadm@sbr-blr-vm3:~>perl AddPoolRangeBatch.pl
```

Usage: **AddPoolRangeBatch.pl** **[-h]** **[-i]** **[-y]** **[-p#]** **[-s]** **<batch file>**, where **<batch file>** could be **<PoolName>** **<Range Start IP>** **<Range End IP | count>**.

The other variables used with the **AddPoolRangeBatch.pl** script are described as follows:

- [-h]—Displays help.
- [-i]—Ignores range conflicts.
- [-y]—Answers yes to questions.
- [-p#]—Uses # processes simultaneously for batch processing, where # represents the number of processes used.
- [-s]—Installs the stored procedure used for cleaning the zombie pool (that is assigning IP from zombie pools to other pools).
- [-u]—Cleans zombie sessions.

**NOTE:** You can alternatively use the **AddRange.sh** script to manage IP pools but it takes longer time to accomplish the task. For more information on **AddRange.sh**, see the *Administration and Configuration Guide*.

# When and How to Restart Session State Register Nodes, Hosts, and Clusters

## IN THIS CHAPTER

- Overview of Starting and Stopping a Session State Register Cluster | 196
- Proper Order for Starting Nodes in a Cluster | 201
- Proper Order for Stopping Individual Nodes | 203
- Stopping a Single Node | 204
- Starting a Single Node | 204
- Proper Order for Stopping the Entire Cluster | 205

This chapter explains how to stop and restart Session State Register nodes or the entire cluster. These topics are in this chapter:

## Overview of Starting and Stopping a Session State Register Cluster

Having to stop all nodes in a cluster is uncommon because most system maintenance can be done on one system at a time. Taking the whole cluster offline defeats the intention of the cluster—to avoid downtime. So ensure that taking all systems down at the same time is required before proceeding. Rather than taking down all nodes, determine whether stopping just the SBR processes or the database management processes might be sufficient.

Stopping a server that hosts both a SBR Carrier and a management node creates a double fault, but does not damage the cluster because a fully redundant cluster always has more than one of each type of node. Stopping multiple nodes that provide redundancy to each other causes multiple faults that may damage the cluster and take the entire cluster off-line.

In the SSR environment, each type of node is started in a specific order so that required resources are available when other nodes require them, and stopped. This means that several commands may be executed on servers that host both SBR Carrier and management nodes.

Startup and shutdown commands must be executed by root on each node.

## sbrd

The **sbrd** script starts and stops processes on Steel-Belted Radius Carrier hosts for all four types of nodes in the cluster. The **sbrd** script may be in either of two directories on servers, depending on whether they have been configured to automatically start all procedures or not using the autoboot functionality which is configured when running the **configure** script.

**NOTE:** The **./configure** script prompts you to enable or disable the autoboot option. If you disable it, you cannot start the SSR process on the node (**./sbrd start ssr**) from the **/etc/init.d/sbrd** directory. If the autoboot option is disabled, you must start the SSR process from the **/opt/JNPR sbr/radius** directory.

All **sbrd** commands are executed by root. In an SSR environment, the **hadm** user can execute the script on SSR processes, but expect errors with RADIUS processes that are owned by root.

### Running sbrd on Session State Register Nodes

This section applies to running **sbrd** on nodes in a Session State Register cluster.

#### Syntax

```
sbrd status [radius|ssr|GWrelay]
sbrd start [radius|ssr|GWrelay] [force]
sbrd start ssr --nowait-nodes=node-ids
sbrd stop [radius|ssr|GWrelay] [force]
sbrd stop [cluster] [force]
sbrd restart [radius|ssr|GWrelay] [force]
sbrd clean [radius|ssr] [force]
sbrd hup [radius|ssr|authGateway [process-name]]
sbrd status [radius|ssr|GWrelay] -v [-p <LCI password>]
```

#### Options

- The **start**, **stop**, and **restart** arguments start, stop, and restart the process. If a subsystem is not specified, the command works only on RADIUS and GWrelay processes because SSR processes normally are not stopped; to stop them, **ssr** must be invoked. For example: **sbrd stop ssr**.
- Executing **stop cluster** on a SBR Carrier server stops both SSR and RADIUS processes. Executing **stop cluster** on a management node also stops the data nodes controlled by the management node.
- The **clean** argument removes lock files that prevent reinitializing the database more than once. You should use this argument only if things go wrong during the initial installation and configuration.

When it is executed on a data node, **clean** also prepares the node to take part in a new environment; for example, if an expansion kit is added to increase the number of data nodes from two to four.

- The **radius**, **ssr**, or **GWrelay** optional argument specifies which process to operate on when executed on a server that hosts more than one node.
  - **radius** specifies the local Steel-Belted Radius Carrier processes
  - **ssr** specifies data node and management node processes according to the type of node on which it is executed
  - **GWrelay** specifies the GWrelay application.
- Executing **start ssr --nowait-nodes=node-ids** starts the cluster without waiting for the full cluster to be initialized. The *node-ids* variable specifies the comma-separated list of node IDs that are unreachable, for example: **sbrd start ssr --nowait-nodes=51,52**. You must use this argument only if one half of the cluster has network connectivity, but has lost the ability to communicate with the other half. When the network connectivity between the two halves of the cluster is restored, you can start the remaining nodes with the normal startup scripts.
- The **status** option displays information such as SBR package version, SBR process status, and loaded plug-in information.
- The **hup** option operates as the **kill -HUP** command does on SBR Carrier nodes, but does not require the process ID. Executing **sbrd hup authGateway** issues the SIGHUP (1) signal to all the authGateway processes running on SBR Carrier. To issue the SIGHUP (1) signal only to the specific authGateway process, you must execute the hup option with the authGateway process name, for example: **sbrd hup authGateway GMT**.
- The **force** argument makes **sbrd** attempt to disregard or overcome any errors that occur when processing the command. Normal behavior without the argument is to halt on errors. For example, **sbrd start** does not attempt to start software that is already running, but **sbrd start force** ignores a running process. This may produce unintended results, so use **force** with great care.
- The **-v** option displays additional information about the RADIUS process along with basic information such as the SBR package version, SBR process status, and SBR process ID. If you have changed the default Lightweight Directory Access Protocol (LDAP) Configuration Interface (LCI) password, you should use the **-p** option to specify the password. For more information about the RADIUS status information, see [“Displaying RADIUS Status Information” on page 79](#).





**CAUTION:** In the case of a cluster, stopping the RADIUS server does not cause any SSR processes to be stopped. If you want to stop the SSR processes on a SBR/management type node (for example for scheduled maintenance of the machine) then as root, navigate to the **radius/install** subdirectory of the directory in which the **JNPRsbr** package was installed (**/opt/JNPRsbr/radius/install** by default) and:

Execute:

```
./sbrd stop ssr
```

If you want to stop the entire cluster (not usually intended), then on each and every node execute as root:

Execute:

```
./sbrd stop cluster
```

When you stop a cluster, the system prompts you with the following warning:

**WARNING:** This function is capable of stopping multiple nodes. Do not use this function if you intend to stop only one node. Do you intend to stop the entire cluster? (y,n): y Are you sure? (y,n): y Really? (y,n): y

### Examples

This example shows the effect of **sbrd stop ssr** executed on a cluster management node:

```
root@wrx07:~> /opt/JNPRsbr/radius/sbrd stop ssr
Stopping ssr auxiliary processes
Stopping ssr management processes
```

```
Connected to Management Server at: 172.28.84.36:5235
Node 1 has shutdown.
Disconnecting to allow Management Server to shutdown
```

This example shows the effect of **sbrd start ssr** on a management node. Be aware that this does not start the data nodes.

```
root@wrx07:~> /opt/JNPRsbr/radius/sbrd start ssr
Starting ssr management processes
bash-3.00#
```

### Notes

- When **sbrd** is executed without a **<radius|ssr>** argument, it runs against all node processes on the server. For example, **sbrd start** starts both RADIUS and SSR processes for all nodes on a server. For complete details see [“When and How to Restart Session State Register Nodes, Hosts, and Clusters” on page 196](#).
- In an SSR environment, because some servers may host both SBR Carrier and management nodes, **sbrd** may be executed more than once with different arguments.
- The **clean** argument removes lock files that prevent reinitializing the database more than once. You should use this argument only if something goes wrong during the initial installation and configuration, or when adding data nodes.

### Starting the Cluster

If all nodes in the cluster are shut down, restarting requires bringing each type of node online in a specific order. If the systems are completely shut down, rebooting the machine restarts the appropriate processes automatically because automatic restart is the default configuration for all types of Session State Register nodes.

If the systems have not been totally shut down and just the SSR processes have been halted, log in as root and execute the **start** commands in the order described in [“Proper Order for Starting Nodes in a Cluster” on page 201](#), to start each type of node’s processes.

During the cluster startup process, each time a SSR or RADIUS process is started on a node, we recommend that you verify the status of that node before moving on to the next node by executing the **sbrd status** command:

1. Log in to the node as hadm or root.
2. Execute:

```
/opt/JNPRsbr/radius/sbrd status
```

Results similar to this example are displayed:

```
hadmUser$>/opt/JNPRsbr/radius/sbrd status
```

```
[ndbd(NDB)]      2 node(s)
id=10    @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11    @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)]  2 node(s)
id=1      @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2      @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
```

```
[mysqld(API)]    4 node(s)
id=21    @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=22    @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=30    @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=31    @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
```

Examine the line starting with **id=**, and verify that there are no references to **starting**, **connecting**, or **not connected**. Any of these references indicate the process has either not finished starting, or the node is not connected properly. You may need to execute the **sbrd status** command more than once because it only shows a snapshot of activity; the display does not refresh automatically. Do not proceed to the next node until you are sure the process has started properly and the node is connected.

## Proper Order for Starting Nodes in a Cluster

Begin by starting all SSR nodes, one at a time in the following order:

1. *SBR/management nodes (sm):*

On each SBR/management node, one at a time, execute **sbrd start ssr**. Then, check the status of that node by executing **sbrd status**, and ensure the SSR process is running without error. There should be no indication of **starting**, **connecting**, or **disconnected**. Repeat this process until each SBR/management node is started and running without error.

2. *Management nodes (m):*

On each management node, one at a time, execute **sbrd start ssr**. Then, check the status of that node by executing **sbrd status**, and ensure the SSR process is running without error. There should be no

indication of **starting**, **connecting**, or **disconnected**. Repeat this process until each management node is started and running without error.

3. *Data nodes (d):*

On each data node, one at a time, execute **sbrd start ssr**. Then, check the status of that node by executing **sbrd status**, and ensure the SSR process is running without error. There should be no indication of **starting**, **connecting**, or **disconnected**. Repeat this process until each data node is started and running without error.

4. If this is the *first time* starting the cluster, or if the *schema has changed*, you need to execute **CreateDB.sh** on every management node before proceeding to the next step of starting the RADIUS process on all SBR nodes.

By definition, the schema has changed if **DestroyDB.sh** has been executed on any management node including: SBR/management nodes (sm) or management nodes (m).

If you have executed **DestroyDB.sh** on any management node, we recommend you run **./sbrd clean** before starting the SSR process on the node with the **./sbrd start ssr** command.

**NOTE:** Except when migrating from a temporary cluster, all SSR processes must be up on all SSR nodes [sm, m, d] and all SBR processes must be down on all SBR nodes [s, sm] in order to execute **CreateDB.sh**. This is also the case when executing **DestroyDB.sh**. See [Table 34 on page 202](#) for a definition of SSR nodes compared to SBR nodes.

**Table 34: Node Type Definitions**

SSR Nodes	SBR Nodes
SBR/management (sm)*	SBR/management (sm)
Management nodes (m)	SBR node (s)
Data nodes (d)	–

\* These are also referred to as *management nodes*.

If this is not the first time starting the cluster, and you are sure the *schema has not changed*, proceed to next Step.

5. Start all SBR nodes, one at a time, by executing **sbrd start radius** in the following order:

a. *SBR nodes (s):*

On each SBR node, one at a time, execute **sbrd start radius**. Then, check the status of that node by executing **sbrd status**, and ensure the RADIUS process is running without error. There should be no indication of **starting**, **connecting**, or **disconnected**. Repeat this process until each SBR node is started and running without error.

b. *SBR/management nodes (sm):*

On each SBR/management node, one at a time, execute **sbrd start radius**. Then, check the status of that node by executing **sbrd status**, and ensure the RADIUS process is running without error. There should be no indication of **starting**, **connecting**, or **disconnected**. Repeat this process until each SBR node is started and running without error.

**NOTE:** When CCM is enabled, you must start the primary SBR Carrier server first, wherever it is. In most cases, one of the sm nodes is the primary. So, in the case of CCM, the order in which SBR nodes are started is ultimately determined by which s or sm node is the primary. Start the SBR nodes beginning with the primary first, and then follow the order listed previously.

## Proper Order for Stopping Individual Nodes

Follow this procedure in precise order to properly shut down nodes in the cluster. Log in as root on each machine, in the order specified, and execute the specified command.



**CAUTION:** Stopping multiple systems and processes removes all redundancy, can create multiple faults, and may damage the cluster. If you do need to stop multiple nodes, be sure to restart them properly. See [“Starting the Cluster” on page 200](#) for the correct order.

1. On each *SBR node (s)*, one at a time, execute:  
**sbrd stop radius**  
**sbrd status**
2. On each *SBR/management (sm) node*, one at a time, execute:  
**sbrd stop radius**  
**sbrd status**
3. On each *SBR/management (sm) node*, one at a time, execute:  
**sbrd stop ssr**  
**sbrd status**
4. On each *management node (m)*, one at a time, execute:  
**sbrd stop ssr**  
**sbrd status**
5. On each *data node (d)*, one at a time, execute:  
**sbrd stop ssr**

**sbrd status**

## Stopping a Single Node

You can stop any single node in the cluster to perform maintenance without affecting the integrity of the cluster (because the server's redundant partner takes on the primary role). These commands just stop the node's process—they have no effect on the system itself, which may still need to be shut down.

If the node is a SBR node, modifying configuration files often requires a restart.

The stop commands for each type of node are as follows:

*SBR nodes (s):*

**sbrd stop radius**

*SBR/management nodes (sm) (in the following order):*

**sbrd stop radius**

**sbrd stop ssr**

*Management nodes (m) and data nodes (d):*

**sbrd stop ssr**

## Starting a Single Node

To restart a node use the following commands.

*SBR nodes (s):*

**sbrd start radius**

**sbrd status** (Verify the node is working error free.)

*SBR/management nodes (sm) (in the following order):*

**sbrd start ssr**

**sbrd status** (Verify the node is working error free.)

**sbrd start radius**

**sbrd status** (Verify the node is working error free.)

*Management nodes (m) and data nodes (d):*

**sbrd start ssr**

**sbrd status** (Verify the node is working error free.)

## Proper Order for Stopping the Entire Cluster

If you intend to shut down the entire cluster, use the following procedure.



**CAUTION:** Before proceeding with this procedure, be sure that you want to shut down all nodes in the cluster.

When you stop a cluster, the system prompts you with the following warning:

**WARNING:** This function is capable of stopping multiple nodes. Do not use this function if you intend to stop only one node. Do you intend to stop the entire cluster? (y,n): y Are you sure? (y,n): y Really? (y,n): y



**CAUTION:** Stopping multiple systems and processes removes all redundancy, can create multiple faults, and may damage the cluster. If you do need to stop the cluster, be sure to restart the cluster properly. See [“Starting the Cluster” on page 200](#).

Follow this procedure in precise order to properly shut down the cluster. Log in as root on *each machine*, in the order specified, and execute the specified command.

1. On each *SBR node (s)*, one at a time, execute:  
**sbrd stop cluster**  
**sbrd status**
2. On each *SBR/management node (sm)*, one at a time, execute:  
**sbrd stop cluster**  
**sbrd status**
3. On each *management node (m)*, one at a time, execute:  
**sbrd stop cluster**  
**sbrd status**
4. On each *data node (d)*, one at a time, execute:  
**sbrd stop cluster**  
**sbrd status**

# 5

PART

## Expanding a Session State Register Cluster

---

Expanding an SSR Cluster | **207**

---



# Expanding an SSR Cluster

## IN THIS CHAPTER

- [Upgrading from a Restricted Cluster to a Regular Cluster | 207](#)
- [Adding Nodes to a Cluster Overview | 210](#)
- [Adding a New SBR Carrier Server to an Existing Cluster | 211](#)
- [Adding a Management Node Expansion Kit to an Existing Cluster | 235](#)
- [Adding a Data Expansion Kit to an Existing Cluster | 259](#)

This chapter explains how to scale an existing Session State Register (SSR) cluster by installing a Data Expansion Kit that adds two data nodes, a third management node, or additional SBR Carrier front end servers to a cluster.

These topics are in the chapter:

## Upgrading from a Restricted Cluster to a Regular Cluster

This section describes how to upgrade from a restricted cluster to a regular cluster. You upgrade from a restricted cluster to a regular cluster to remove the restrictions on the maximum number of concurrent sessions and you are allowed to add an expansion kit to your cluster. To upgrade to a regular cluster, perform the following tasks:

**NOTE:** This procedure is applicable only if you provided a Restricted Cluster license during SSR installation.

This procedure must be executed in all SBR nodes (that is, S and SM nodes) involved in the SSR cluster.

1. As root, navigate to the directory where you installed the Steel-Belted Radius Carrier package. For information about directory in which Steel-Belted Radius Carrier package is installed, see [“Unpacking Session State Register Software” on page 95](#).

Then, navigate to the **radius/install** subdirectory and run:

```
cd /opt/JNPRsbr/radius/install/
```

2. Run the **configure** script:

```
./configure
```

3. At the configuration tasks prompt, enter **6** to upgrade from a restricted cluster to a regular cluster.

Configuring SBR Software

```
-----
SBR 8.60.24028 cluster red{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-2.englab.juniper.net(sm)
is CONFIGURED and processes are UP, may be stopped if reconfigured
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on any node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

Enter number of desired configuration task, or q to quit : **6**

#### 4. Enter the license number of the Regular Cluster license.

Enter the Upgrade regular cluster license:

Restricted Cluster has been upgraded to regular cluster.

Now as per your concurrent session license, maximum concurrent sessions will bring into effect.

For adding concurrent session license, use the SBR Admin Web GUI.

#### 5. Execute the steps 1 to 4 in all SBR nodes (that is, S and SM nodes) involved in the SSR cluster.

After you have successfully upgraded from a restricted cluster to a regular cluster, the restriction on the maximum number of concurrent sessions is removed and you are allowed to add any number of concurrent session licenses through the Web GUI.

## Adding Nodes to a Cluster Overview

Adding any type of node to an existing cluster is very similar to the initial installation. Adding SBR nodes or a management node is easier than adding data nodes because the first two types of nodes do not directly support the database. Adding data nodes requires taking the existing data nodes offline and rebuilding the session database.

In general, to add new nodes to an existing cluster:

1. Plan the integration of the new nodes by filling out an expansion node worksheet. See [Table 10 on page 35](#).
2. Work through the tasks and prerequisites in [“Preparing for a Steel-Belted Radius Carrier Installation” on page 31](#) to ensure that the new servers meet all SSR server requirements.
3. On a management node in the existing cluster, build upon the existing cluster configuration to generate updated cluster definition files that outline the new cluster topology.

Use the entries from [Table 10 on page 35](#) to provide the names, IP addresses, and node IDs of the new servers.

4. Transfer a copy of the **tar** file containing the new cluster definition files to each existing node and each new node.
5. Install SBR Carrier software on the new nodes, and configure the existing cluster nodes with new topology.
6. Restart the cluster so all nodes come online using the new cluster definition files.

The difference between adding SBR Carrier nodes or a management node and data nodes manifests itself in the way Steps 5 and 6 are implemented.

- When installing SBR Carrier nodes or a management node, each node is reconfigured and restarted, one at a time. Because only one node at a time is out of service, the cluster’s redundancy protects all connections.
- Adding data nodes requires taking the existing data nodes offline and rebuilding the SSR database, so all nodes are reconfigured, then all nodes are restarted as a cluster.

**NOTE:** The procedures in this chapter describe how to add an SBR node, a management node, or data nodes to an existing cluster. These procedures assume a basic cluster configuration of two (sm) nodes and two (d) nodes. Because your cluster configuration may differ from this, review these procedures in their entirety before starting the procedure. If your configuration is different, you may need to make changes to these procedures. If you have to modify these procedures, see [“When and How to Restart Session State Register Nodes, Hosts, and Clusters” on page 196](#), to determine the proper order for starting and stopping the various node types in your cluster.

## Adding a New SBR Carrier Server to an Existing Cluster

This section describes how to add additional SBR Carrier servers to an existing cluster. To add an SBR node to an existing cluster, you perform the following high-level tasks:

1. Update the existing cluster definition files.  
See [“Updating the Existing Cluster Definition Files for the New SBR Node” on page 212.](#)
2. Distribute the updated cluster definition files to the existing nodes in the cluster.  
See [“Distributing the Updated Cluster Definition Files to the Existing Nodes” on page 218.](#)
3. Install the SBR Carrier software on the new SBR node.  
See [“Installing the SBR Carrier Software on the New SBR Node Host Machine” on page 220.](#)
4. Configure the SBR Carrier software on the new SBR node.  
See [“Configuring the Software on the New SBR Node” on page 222.](#)
5. One by one, stop the process on each existing node, configure it with the new cluster definition file, and restart the process.  
See [“Configuring Each Existing Node in the Cluster with the New Cluster Definition Files” on page 228.](#)
6. Start the RADIUS process on the new SBR node.  
See [“Starting the New SBR Node” on page 235.](#)

The following procedure adds a single SBR node to the existing cluster. The procedure is the same when adding multiple SBR nodes to an existing cluster.

The following designations are used throughout the examples in this section:

sm = Hardware has SBR node and Management node.

s = Hardware has only SBR node.

m = Hardware has only Management node.

d = Hardware has Data node.

2sm, 2d = Two SBR/Management nodes and 2 Data nodes.

2S, 2SM, 2D = Two SBR nodes and 2 SBR/Management nodes, 2 Data nodes.

Display the existing cluster:

```
hadm@wrx07:~> ndb_mgm -e show
Connected to Management Server at: 172.28.84.166:5235
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
```

```
id=10    @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11    @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1      @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2      @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
```

```
[mysqld(API)] 4 node(s)
id=6      @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=7      @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=58     @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=59     @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
```

The existing cluster includes two full SBR Carrier licenses and a license for the Starter Kit resulting in a configuration that includes two sm nodes and two d nodes.

For the purposes of this procedure, the existing two sm nodes are identified as sm1 and sm2 as follows:

```
id=1 @172.28.84.36 = sm1
id=2 @172.28.84.166 = sm2
```

## Updating the Existing Cluster Definition Files for the New SBR Node

In this first part of the procedure, you update the existing cluster definition files on SM1 to reflect the new configuration: 1s, 2sm, 0m, 2d.

Before proceeding, make sure the machine that you want to host the new SBR node meets all system requirements. See [“Before You Install Software” on page 39](#).

The following steps create a new set of cluster definition files in `/opt/JNPRshare/install/ <cluster_name>` and in `configure.<cluster_name>.tar`. You may want to make a backup copy of the existing `configure.<cluster_name>.tar` file before creating the new files, in case you need to restore the existing configuration.

To generate the updated cluster definition files:

1. As root, on the sm1 node, navigate to the **radius/install** subdirectory of the directory in which the **JNPRsbr** package was installed (**/opt/JNPRsbr** by default).

Example:

```
cd /opt/JNPRsbr/radius/install/
```

2. Run the **configure** script:

Execute:

```
./configure
```

Example:

```
root@sbrha-4:/opt/JNPRsbr/radius/install> ./configure
```

```
Configuring SBR Software
```

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is CONFIGURED and processes are UP, may be stopped if reconfigured
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (4,q):

### 3. From the menu of configuration tasks, enter **2** to specify **Generate Cluster Definition**.

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is CONFIGURED and processes are UP, may be stopped if reconfigured
-----
```

Generating Cluster Definition...

Enter SBR cluster name [cambridge]:

You are prompted to enter the name of the cluster.

### 4. Press Enter to use the current cluster name.



You are prompted either to create a new cluster or update an existing cluster definition.

```
Create (c) new or update (u) existing cluster definition? [u]:
```

5. Enter **u** to update the existing cluster definition.

The SBR Cluster Starter Kit license allows you to create a minimal cluster of 2 SBR nodes, 2 management nodes, and 2 data nodes. When each node is installed on a separate machine the cluster topology is denoted as {2s,2m,2d}. When SBR nodes are paired with management nodes on the same machines the cluster topology is denoted as {2sm,2d}.

An optional SBR Cluster Management Expansion Kit allows you to add a third management node for {2sm,1m,2d} and an optional Data Expansion Kit allows you to add 2 more data nodes for {2sm,1m,4d} clusters. Additional SBR licenses allow you to add up to 18 more SBR nodes to obtain a maximal cluster {18s,2sm,1m,4d} and/or enable extra features.

While it is not difficult to add management and/or SBR nodes to an existing cluster, adding data nodes is more difficult and may require you to shutdown the entire cluster as opposed to a rolling restart.

Another license is required if you wish to add a third management node. Adding a third management node will require a rolling restart later.

Enter Management Expansion Kit license, if any:  
 Another license is required if you wish to add more data nodes.  
 Adding data nodes may require you to shutdown the entire cluster.  
 Enter Data Expansion Kit license, if any:  
 This cluster presently contains 2 of 20 possible SBR nodes.  
 Adding more SBR nodes will require a rolling restart later.  
 Enter number of SBR nodes to be added [0]:

6. Enter **1** to add one SBR node.

```
Updating cluster cambridge{1s,2sm,0m,2d}
will require 1 new machines. Do you wish to continue? [y]:
```

7. Verify the proper configuration of **{1s,2sm,0m,2d}** for the cluster named **cambridge** and enter **y** to continue.

```
Information will now be gathered for each new machine to be added.
You will have a chance to review all information at least once
before any machines are modified.
```

8. When prompted, enter the IP address and license number for the new SBR node.

```
-----
SBR 8.60.50006 cluster cambridge{1s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
Partial configuration at present is {0s,2sm,0m,2d} of {1s,2sm,0m,2d}
-----
IMPORTANT: node names must be entered as reported by 'uname -n'.
Enter node name [cambridge-5]: sbrha-8.carrier.spgma.juniper.net
Enter node type (s) [s]:
Enter SBR node ID (40-59) [57]:
Enter SBR node IP address by which it is known to management nodes.
Enter SBR node IP address: 172.28.84.178
Enter SBR licenses meant only for this particular SBR node.
Enter one license per line and an empty line when finished.
Enter SBR full license: xxxx xxxx xxxx xxxx
Enter SBR feature license:
```

```
-----
SBR 8.60.50006 cluster cambridge{1s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
Generated configuration is {1s,2sm,0m,2d} of {1s,2sm,0m,2d}
-----
```

```
Generating configuration files
```

The system generates the required cluster definition files and prompts you to view, accept, or reject them.

9. Enter **a** to accept them and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files: a
```

```
Writing shared configuration to /opt/JNPRshare/install/cambridge
```

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is CONFIGURED and processes are UP, may be stopped if reconfigured
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

READY: last operation succeeded, generated cluster definition.  
Enter the number of the desired configuration task or quit (4,q):

The software writes the new cluster definition files to this node and returns you to the main configuration menu.

10. Press **q** to quit.

## Distributing the Updated Cluster Definition Files to the Existing Nodes

At this point, the updated cluster definition files (\*.tar file) have been created on the sm1 node only. Now you need to distribute the new definition files to the other nodes in the cluster, including the new SBR node.

On both the existing nodes and the new SBR node in the expanded cluster, create a copy of the new cluster definition files. Doing this does not invoke the new files, but makes them available to the **configure** script later in the workflow.

To distribute the new cluster definition files:

1. Log in to each node (existing and new) as **hadm**.
2. Change directories to the install directory.

(On new nodes, the entire path may not exist because the **<cluster name>** portion of the path was not created when you prepared the new machine, so you may need to create it.) See [“Creating Share Directories” on page 52](#).

Execute:

```
cd /opt/JNPRshare/install/ <cluster_name>
```

For example:

```
cd /opt/JNPRshare/install/cambridge
```

3. Use FTP binary mode to connect to the node host (in this example sm1) where you created the new cluster definition files.
4. Execute the **get** command to transfer the **configure. <cluster name> .tar** file.

For example:

```
bin
```

```
get /opt/JNPRsbr/radius/install/configure.cambridge.tar
```

5. In a terminal window, extract the new cluster definition files from the archive.

Execute:

```
tar xvf configure. <cluster_name> .tar
```

Output similar to this example is displayed:

```
$ tar xvf configure.MyCluster.tar
x dbcluster.rc, 1925 bytes, 4 tape blocks
x config.ini, 2435 bytes, 5 tape blocks
x my.cnf, 1017 bytes, 2 tape blocks
x dbclusterndb.gen, 33474 bytes, 66 tape blocks
x dbcluster.dat, 921 bytes, 2 tape blocks
```

6. Repeat these steps until every node in the cluster has a copy of the new cluster definition files.

## Installing the SBR Carrier Software on the New SBR Node Host Machine

This procedure describes how to unpack and install the SBR Carrier software on the host machine for the new SBR node.

1. Log in to the machine as root.
2. Copy the Steel-Belted Radius Carrier installation files from their download location to the machine. Make sure to copy them to a local or remote hard disk partition that is readable by root.

This example copies the files from a download directory to the **/tmp/sbr** directory.

Execute:

```
mkdir -p /opt/tmp  
cp -pR /tmp/sbr/solaris/* /opt/tmp/
```

3. Extract the SBR Carrier installation package.

For 64-bit Solaris, execute:

```
cd /tmp/sbr  
ls -ltr
```

```
total 216240  
-rw-r--r--  1 root      root      110712276 Aug 25 09:44  
sbr-cl-8.6.0.R-1.sparcv9.tgz
```

Execute:

```
gunzip sbr-cl-8.6.0.R-1.sparcv9.tgz  
tar xf sbr-cl-8.6.0.R-1.sparcv9.tar
```

4. Verify that the extraction worked and confirm the name of the package file.

For 64-bit Solaris, execute:

```
ls -ltr
```

```
total 216256  
drwxr-xr-x  4 Xtreece  other      370 Aug 24 17:01 JNPRsbr.pkg  
-rw-r--r--  1 root      root      110712276 Aug 25 09:44  
sbr-cl-8.6.0.R-1.sparcv9.tar
```

5. Install the package.

Execute:

```
pkgadd -d /tmp/sbr
```

```
The following packages are available:
```

```
1  JNPRsbr.pkg          JNPRsbr - Juniper Networks Steel-Belted Radius
    (Carrier Cluster Edition)          (sparc) 8.60.50006
```

```
Select package(s) you wish to process (or 'all' to process all packages).
(default: all) [?,??,q]:
```

6. Type **all** and press Enter.

The script resumes.

```
Processing package instance <JNPRsbr.pkg> from </tmp>
```

7. Confirm the installation directory.

Depending on the system configuration, you are prompted whether to create the **/opt/JNPRsbr** directory if it does not exist, or to over-write an already extracted package, or any of several other questions.

```
The selected base directory </opt/JNPRsbr> must exist before installation is
attempted.
```

```
Do you want this directory created now [y,n,?,q]
```

8. Answer the question appropriately (or change the extraction path if necessary) so that the script can proceed.

To accept the default directory as a target, enter **y**.

The script resumes.

```
Using </opt/JNPRsbr> as the package base directory.
#Processing package information.
#Processing system information.
    48 package pathnames are already properly installed.
#Verifying disk space requirements.
```

```
#Checking for conflicts with packages already installed.
#Checking for setuid/setgid programs.
```

```
This package contains scripts which will be executed with super-user
permission during the process of installing this package.
```

```
Do you want to continue with the installation of <JNPRsbr> [y,n,?]
```

9. Enter **y** to confirm that you want to continue to install the package.

```
Installing JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster
Edition) as <JNPRsbr>
```

```
## Executing preinstall script.
## Installing part 1 of 1.
.
.
.
[ verifying class <none> ]
## Executing postinstall script.
Newly installed server directory will be backed up as:
/opt/JNPRsbr/radius/install/backups/2009:03:31-00:34:06
```

```
Installation of <JNPRsbr> was successful.
```

## Configuring the Software on the New SBR Node

Before starting this procedure, review [“Before You Install Software” on page 39](#). In particular, review requirements for [“Setting Up External Database Connectivity \(Optional\)” on page 53](#) and [“Installing the SIGTRAN Interface \(Optional\)” on page 55](#), because steps in this procedure require the server to be preconfigured for these capabilities.



To install the software on the new SBR node:

1. As root, navigate to the **radius/install** subdirectory of directory where you installed the Steel-Belted Radius Carrier package in [“Installing the SBR Carrier Software on the New SBR Node Host Machine” on page 220](#).

Example: **cd /opt/JNPRsbr/radius/install**

2. Run the **configure** script.

Execute:

**./configure**

3. Review and accept the Steel-Belted Radius Carrier license agreement.

Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

4. From the menu of configuration tasks, enter **3** to specify **Configure Cluster Node**.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-8.carrier.spgma.juniper.net
is not configured and processes are down, needs to be configured
-----
```

```
1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.
```

```
2.  Generate Cluster Definition
    Creates new or updates existing cluster definitions.
    Modifies the shared directory but does not modify this node.
```

```
3.  Configure Cluster Node
    To be preceded by 'Generate Cluster Definition' on one node.
    Must be invoked on each and every node of the cluster.
```

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q): 3

### 5. Specify the name of the cluster.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-8.carrier.spgma.juniper.net
is not configured and processes are down, needs to be configured
-----
```

Configuring Cluster Node...

Enter SBR cluster name [sbrha]:cambridge

```
Reading shared configuration from /opt/JNPRshare/install/cambridge
```

```
Generating configuration files
```

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files: a
```

6. Enter **a** to accept the modified configuration files and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

7. The configure script prompts you with a warning about whether or not to apply the cluster definition to this node. Enter **y** to proceed.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y
```

```
Cleaning directories
/opt/JNPRhadm
```

```
Applying configuration
```

```
Initializing Session State Register, please wait a few minutes...
```

**NOTE:** Expect package is required to be installed on RHEL 7.3 or later (see [“Linux” on page 47](#) for supported RHEL versions). SBR Carrier has been tested only with the Expect-5.45 version. Expect package is installed on Solaris 11.3.36.10.0 or later version by default (see [“Solaris” on page 47](#) for supported Solaris versions).

## 8. Configure the SBR node.

For information about responding to the following prompts, see [“Configuring the Host Software on the First Server in the Cluster” on page 99](#).

```
Do you want to configure Java Runtime Environment for JDBC Feature [n] :
Please enter backup or radius directory from which to migrate.
Enter n for new configuration, s to search, or q to quit
[n]:
```

```
Enter initial admin user (UNIX account must have a valid password) [root]:
Enable Centralized Configuration Management (CCM) for this SBR node? [n]:
Configuring for use with generic database
Do you want to enable "Radius WatchDog" Process? [n]:
Do you want to enable LCI? [n]:
Do you want to configure for use with Oracle? [n]:
Removing oracle references from startup script
Do you want to configure for use with SIGTRAN? [n]:
Removing SIGTRAN references from startup script
Do you want to configure SNMP? [n]:
Configuring Admin GUI Webserver
Compatible Java version found :
Do you want to install custom SSL certificate for Admin WebServer? [n]:
Enable (e), disable (d), or preserve (p) autoboot scripts [e]:
```

```
The SBR Admin Web GUI can be launched using the following URL:
https://<servername>:2909
```

```
Configuration complete
```

```
-----
SBR 8.60.50006 cluster cambridge{1s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-8.carrier.spgma.juniper.net(s)
is configured and processes are down, may be reconfigured if desired
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

```
READY: last operation succeeded, node configured.
Enter the number of the desired configuration task or quit (4,q):
```

9. Enter **q** to quit.

### Configuring Each Existing Node in the Cluster with the New Cluster Definition Files

At this point in the process, all nodes in the cluster have the new cluster definition files loaded and the new SBR node has been configured with the new files. The existing nodes are still running with the old cluster definition files.

In this procedure, you log in to each existing node, stop the process, run the **configure** script, and restart the processes. You must complete these steps on every existing node in the cluster. This example starts with the sm1 node.



**CAUTION:** In this procedure, you need to stop and restart each existing node one by one to apply the new cluster definition to each of the original cluster nodes and to *keep the cluster up and running*. Do not operate on multiple nodes at the same time because that creates multiple faults that can stop the entire cluster.

Always review the recommended start and stop order and processes and plan the order in which to perform the equivalent steps in your cluster. See [“When and How to Restart Session State Register Nodes, Hosts, and Clusters” on page 196.](#)

1. Log in to the first existing node (in this example, sm1) as root.
2. Navigate to the radius subdirectory of the directory in which the **JNPRsbr** package was installed (/opt/JNPRsbr by default).

Example: `cd /opt/JNPRsbr/radius`

3. Stop the RADIUS processes on the node you are configuring:
  - a. Execute:

```
./sbrd stop radius
```

b. Execute:

```
./sbrd status
```

4. Stop the **ssr** processes on the node you are configuring:

```
./sbrd stop ssr
```

5. Check the status of the node:

```
./sbrd status
```

```
Cluster Configuration
```

```
-----
```

```
[ndbd(NDB)]      2 node(s)
```

```
id=10  @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
```

```
id=11  @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)]  2 node(s)
```

```
id=1 (not connected, accepting connect from 172.28.84.36)
```

```
id=2   @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)
```

```
[mysqld(API)]    4 node(s)
```

```
id=6 (not connected, accepting connect from 172.28.84.36)
```

```
id=7   @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)
```

```
id=58  @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)
```

```
id=59 (not connected, accepting connect from 172.28.84.36)
```

```
hadm@wrx07:~>
```

6. Verify that the node you are about to configure is not connected. In this example, the node ID for sm1 indicates: **id=1 (not connected, accepting connect from 172.28.84.36)**, indicating the sm1 node is stopped.

7. Run the **configure** script:

Execute:

```
./configure
```

## Configuring SBR Software

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is configured and processes are down, may be reconfigured if desired
-----
```

### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.



#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (4,q):

8. From the menu of configuration tasks, enter **3** to specify **Configure Cluster Node**.

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is configured and processes are down, may be reconfigured if desired
-----
```

Configuring Cluster Node...

Enter SBR cluster name [cambridge]:

9. Press Enter to accept the existing cluster name (in the example: cambridge).

Create (c) new or update (u) existing node configuration? [u]:

10. Enter **u** to update the existing node configuration.

Reading shared configuration from /opt/JNPRshare/install/cambridge

Generating configuration files

```

Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files:

```

11. Enter **a** to accept the new cluster definition files.

```

WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n):

```

The configure script prompts you with a warning about whether or not to apply the new cluster definition to the node.

12. Enter **y** to continue.

```

Applying configuration

```

```

-----
SBR 8.60.50006 cluster cambridge{1s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is configured and processes are down, may be reconfigured if desired
-----

```

```

1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.

```

## 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.

## 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

## 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

## 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

## 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

READY: last operation succeeded, node configured.  
Enter the number of the desired configuration task or quit (4,q):

13. Enter **q** to quit.

```
root@sbrha-4:/opt/JNPRsbr/radius/install>
```

14. Notice that the first line in the applied configuration is: **SBR 8.60.50006 cluster cambridge{1s,2sm,0m,2d}**, indicating that the new configuration has been applied to the sm1 node.

15. Restart the SSR process on the newly configured sm1 node:

- a. Execute:  
**./sbrd start ssr**
- b. Execute:  
**./sbrd status**
- c. Examine each line and ensure the SSR process is running without error.

16. Restart the RADIUS process on sm1:

- a. Execute:  
**./sbrd start radius**
- b. Execute:  
**./sbrd status**

```
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
id=10   @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11   @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1    @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2    @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
```

```
[mysqld(API)]   4 node(s)
id=6    @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=7    @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=58   @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=59   @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
```

- c. Examine each line and ensure the RADIUS process is running without error.

17. Repeat steps 1 through 16 on sm2 node and then on each data node, one at a time. For the data nodes, you do not need to stop and restart the RADIUS process because data nodes only run the SSR process.

Do not operate on multiple nodes at once. Doing so creates multiple faults that can stop the entire cluster. Complete the procedure on each node one at a time until the node is operating without error.

## Starting the New SBR Node

At this point in the process, the original nodes in the cluster (sm1, sm2, d1, and d2) are up and running with the new cluster definition files and the new SBR node is configured with the proper configuration, but is not yet running in the cluster. This final process starts the RADIUS process on the new SBR node.

1. Start the RADIUS process on the new SBR node:
  - a. Log in as root to the SBR (s) node.
  - b. Change directories to `/opt/JNPRsbr/radius/`.
  - c. Execute:  
`./sbrd start radius`
  - d. Execute:  
`./sbrd status`
  - e. Examine each line of the final cluster configuration and ensure the RADIUS process is running without error.
2. Complete the configuration of the new SBR node using Web GUI. See [“Basic SBR Carrier Node Configuration” on page 138](#). For complete details, see the *SBR Carrier Administration and Configuration Guide*.

## Adding a Management Node Expansion Kit to an Existing Cluster

This section describes how to add a management node to an existing cluster using the Management Node Expansion Kit. The Management Node Expansion Kit provides software and a license for a third management node. This third management node (which is the final arbiter) is set up on a separate host machine as an (m) node. It does not share the machine with an SBR node (sm node). You must place the third management node in a location that has similar connectivity to the database as the bulk of your NAS devices. You must set the ArbitrationRank for the third M node to 1 (ArbitrationRank=1) and the other two M nodes as ArbitrationRank=2. In this case, during an NOC outage, the third management node decides which half of the cluster lives. You may even need to set up a VLAN or a special firewall connection between this M node and the red zone on which your D nodes are networked.

Adding the third management node increases the resiliency of the cluster by providing an additional arbiter in case of a node failure.

To add a new management node to an existing cluster, you perform the following high-level tasks:

1. Update the existing cluster definition files to include the new management node.  
See [“Updating the Existing Cluster Definition Files for the New Management Node” on page 237.](#)
2. Distribute the updated cluster definition files to the existing nodes in the cluster.  
See [“Distributing the Updated Cluster Definition Files to the Existing Nodes” on page 244.](#)
3. Install the SBR Carrier software on the new management node.  
See [“Installing the SBR Carrier Software on the New Management Node Host Machine” on page 245.](#)
4. Configure the SBR Carrier software on the new management node.  
See [“Configuring the SBR Carrier Software on the New Management Node” on page 248.](#)
5. One by one, stop the process on each existing node, configure it with the new cluster definition file, and restart the process.  
See [“Configuring Each Existing Node in the Cluster with the New Cluster Definition Files” on page 252.](#)
6. Start the SSR process on the new management node.  
See [“Starting the New Management Node” on page 259.](#)
7. Run **CreateDB.sh** on the new management node.  
See [“Running CreateDB.sh on the New Management Node” on page 259.](#)

The following procedure adds a single management node to an existing cluster.

The following designations are used throughout the examples in this section:

sm = Hardware has SBR node and Management node.

s = Hardware has only SBR node.

m = Hardware has only Management node.

d = Hardware has Data node.

2sm, 2d = Two SBR/Management nodes and 2 Data nodes.

2S, 2SM, 2D = Two SBR nodes and 2 SBR/Management nodes, 2 Data nodes.

Display the existing cluster:

```
hadm@wrx07:~> ndb_mgm -e show
Connected to Management Server at: 172.28.84.166:5235
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
id=10   @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11   @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1      @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2      @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
```

```
[mysqld(API)] 4 node(s)
id=6      @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=7      @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=58     @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=59     @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
```

The existing cluster includes two full SBR Carrier licenses and a license for the Starter Kit resulting in a configuration that includes: two sm nodes and two d nodes.

For the purposes of this procedure, the existing two sm nodes are identified as sm1 and sm2 as follows:

```
id=1 @172.28.84.36 = sm1
```

```
id=2 @172.28.84.166 = sm2
```

## Updating the Existing Cluster Definition Files for the New Management Node

In this first part of the procedure, you update the existing cluster definition files on SM1 to reflect the new configuration of 0s, 2sm, 1m, 2d.

Before proceeding, make sure the machine that you want to host the new management node meets all system requirements. See [“Before You Install Software” on page 39](#).

The following steps create a new set of cluster definition files in `/opt/JNPRshare/install/ <cluster_name>` and in `configure.<cluster_name>.tar`. You may want to make a backup copy of the existing `configure.<cluster_name>.tar` file before creating the new files, in case you need to restore the existing configuration.

To generate the updated cluster definition files:

1. As root, on the sm1 node, navigate to the **radius/install** subdirectory of the directory in which the JNPRsbr package was installed (**/opt/JNPRsbr** by default).

Example: **cd /opt/JNPRsbr/radius/install**

2. Run the **configure** script.

Execute:

**./configure**

Example:

```
root@wrx07:/opt/JNPRsbr/radius/install> ./configure
Configuring SBR Software
```

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node wrx07(sm)
is CONFIGURED and processes are UP, may be stopped if reconfigured
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.



#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (4,q):

### 3. From the menu of configuration tasks, enter **2** to specify **Generate Cluster Definition**.

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node wrx07(sm)
is CONFIGURED and processes are UP, may be stopped if reconfigured
-----
```

Generating Cluster Definition...

Enter SBR cluster name [cambridge]:

You are prompted to enter the name of the cluster.

### 4. Press Enter to use the current cluster name.

You are prompted either to create a new cluster or update an existing cluster definition.

```
Create (c) new or update (u) existing cluster definition? [u]:
```

5. Enter **u** to update the existing cluster definition.

The SBR Cluster Starter Kit license allows you to create a minimal cluster of 2 SBR nodes, 2 management nodes, and 2 data nodes. When each node is installed on a separate machine the cluster topology is denoted as {2s,2m,2d}. When SBR nodes are paired with management nodes on the same machines the cluster topology is denoted as {2sm,2d}.

An optional SBR Cluster Management Expansion Kit allows you to add a third management node for {2sm,1m,2d} and an optional Data Expansion Kit allows you to add 2 more data nodes for {2sm,1m,4d} clusters. Additional SBR licenses allow you to add up to 18 more SBR nodes to obtain a maximal cluster {18s,2sm,1m,4d} and/or enable extra features.

While it is not difficult to add management and/or SBR nodes to an existing cluster, adding data nodes is more difficult and may require you to shutdown the entire cluster as opposed to a rolling restart.

Another license is required if you wish to add a third management node. Adding a third management node will require a rolling restart later.  
Enter Management Expansion Kit license, if any: 1770 0002 0112 0100 1145 3801

6. Enter the license number for the new management node.

Another license is required if you wish to add more data nodes. Adding data nodes may require you to shutdown the entire cluster.  
Enter Data Expansion Kit license, if any:

7. You are prompted for the license if adding a Data Expansion Kit. Press Enter because we are not adding a Data Expansion Kit.

```
This cluster presently contains 2 of 20 possible SBR nodes.
Adding more SBR nodes will require a rolling restart later.
Enter number of SBR nodes to be added [0]:
```

8. You are prompted for the license if adding an SBR node. Press Enter because we are not adding an SBR node.

```
Updating cluster cambridge{0s,2sm,1m,2d}
will require 1 new machines. Do you wish to continue? [y]:
```

9. Verify the proper configuration of **{0s,2sm,1m,2d}** for the cluster named cambridge and enter y to continue.

```
Information will now be gathered for each new machine to be added.
You will have a chance to review all information at least once
before any machines are modified.
```

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,1m,2d}
on SunOS 5.10 Generic_141444-09 node wrx07(sm)
Partial configuration at present is {0s,2sm,0m,2d} of {0s,2sm,1m,2d}
-----
IMPORTANT: node names must be entered as reported by 'uname -n'.
Enter node name [cambridge-6]: sbrha-8.carrier.spgma.juniper.net
Enter node type (m) [m]:
```

10. Press Enter for the management node.

```
Enter MGMT node ID (1-3) [3]:
```

11. Press Enter to accept the management node ID.

```
Enter MGMT node IP address by which it is known to other nodes.
Enter MGMT node IP address: 172.28.84.178
```

12. Enter the IP address for the new management node and press Enter.

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,1m,2d}
on SunOS 5.10 Generic_141444-09 node wrx07(sm)
Generated configuration is {0s,2sm,1m,2d} of {0s,2sm,1m,2d}
-----
```

```
Generating configuration files
```

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files:
```

The system generates the required cluster definition files and prompts you to view, accept, or reject them.

13. Enter **a** to accept them and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

In this example, notice that the new configuration displays as Generated configuration is **{0s,2sm,1m,2d}**, confirming that the new management node is included in the cluster definition.

Enter **a** to accept the new definition files.

```
Writing shared configuration to /opt/JNPRshare/install/cambridge
```

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
```

```
on SunOS 5.10 Generic_141444-09 node wrx07(sm)
is CONFIGURED and processes are UP, may be stopped if reconfigured
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

```
READY: last operation succeeded, generated cluster definition.
Enter the number of the desired configuration task or quit (4,q):
```

The software writes the new cluster definition files to this node and returns you to the main configuration menu.

14. Press **q** to quit.

## Distributing the Updated Cluster Definition Files to the Existing Nodes

At this point, the updated cluster definition files (\*.tar file) have been created on the sm1 node only. Now you need to distribute the new definition files to the other nodes in the cluster, including the new management node.

On both the existing nodes and the new management node in the expanded cluster, create a copy of the new cluster definition files. Doing this does not invoke the new files, but makes them available to the **configure** script later in the workflow.

To distribute the new cluster definition files:

1. Log in to each node (existing and new) as **hadm**.
2. Change directories to the installation directory.

(On new nodes, the entire path may not exist because the **<cluster name>** portion of the path was not created when you prepared the new machine, so you may need to create it.) See [“Creating Share Directories” on page 52](#).

Execute:

```
cd /opt/JNPRshare/install/ <cluster_name>
```

For example:

```
cd /opt/JNPRshare/install/cambridge
```

3. Use FTP binary mode to connect to the node host (in this example, sm1) where you created the new cluster definition files.
4. Execute the **get** command to transfer the **configure. <cluster name> .tar** file.

For example:

```
bin
```

```
get /opt/JNPRsbr/radius/install/configure.cambridge.tar
```

5. In a terminal window, extract the new cluster definition files from the archive.

Execute:

```
tar xvf configure. <cluster_name> .tar
```

Output similar to this example is displayed:

```
$ tar xvf configure.MyCluster.tar
x dbcluster.rc, 1925 bytes, 4 tape blocks
x config.ini, 2435 bytes, 5 tape blocks
x my.cnf, 1017 bytes, 2 tape blocks
x dbclusterndb.gen, 33474 bytes, 66 tape blocks
x dbcluster.dat, 921 bytes, 2 tape blocks
```

6. Repeat these steps until every node in the cluster has a copy of the new cluster definition files.

## Installing the SBR Carrier Software on the New Management Node Host Machine

This procedure describes how to unpack and install the SBR Carrier software on the host machine for the new management node.

1. Log in to the host machine for the new management node as root.
2. Copy the Steel-Belted Radius Carrier installation files from their download location to the machine. Make sure to copy them to a local or remote hard disk partition that is readable by root.

This example copies the files from a download directory to the **/tmp/sbr** directory.

Execute:

```
mkdir -p /opt/tmp
cp -pR /tmp/sbr/solaris/* /opt/tmp/
```

3. Extract the SBR Carrier installation package.

For 64-bit Solaris, execute:

```
cd /tmp/sbr
ls -ltr
```

```
total 216240
-rw-r--r--  1 root      root      110712276 Aug 25 09:44
sbr-cl-8.6.0.R-1.sparcv9.tgz
```

Execute:

```
gunzip sbr-cl-8.6.0.R-1.sparcv9.tgz
tar xf sbr-cl-8.6.0.R-1.sparcv9.tar
```

4. Verify that the extraction worked and confirm the name of the package file.

For 64-bit Solaris, execute:

```
ls -ltr
```

```
total 216256
drwxr-xr-x   4 Xtreece  other      370 Aug 24 17:01 JNPRsbr.pkg
-rw-r--r--   1 root     root       110712276 Aug 25 09:44
sbr-cl-8.6.0.R-1.sparcv9.tar
```

5. Install the package.

Execute:

**pkgadd -d /tmp/sbr**

```
The following packages are available:
1  JNPRsbr.pkg          JNPRsbr - Juniper Networks Steel-Belted Radius
   (Carrier Cluster Edition)          (sparc) 8.60.50006
```

```
Select package(s) you wish to process (or 'all' to process all packages).
(default: all) [?,??,q]: all
```

6. Type **all** and press Enter.

The script resumes.

```
Processing package instance <JNPRsbr.pkg> from </tmp>
```

7. Confirm the installation directory.

Depending on the system configuration, the script prompts you to create the **/opt/JNPRsbr** directory if it does not exist, or to over-write an already extracted package, or any of several other questions.

```
The selected base directory </opt/JNPRsbr> must exist before installation is
attempted.
```

```
Do you want this directory created now [y,n,?,q]
```

8. Answer the question appropriately (or change the extraction path if necessary) so that the script can proceed.

To accept the default directory as a target, enter **y**.

The script resumes.



```

Using </opt/JNPRsbr> as the package base directory.
#Processing package information.
#Processing system information.
    48 package pathnames are already properly installed.
#Verifying disk space requirements.
#Checking for conflicts with packages already installed.
#Checking for setuid/setgid programs.

```

This package contains scripts which will be executed with super-user permission during the process of installing this package.

Do you want to continue with the installation of <JNPRsbr> [y,n,?]

9. Enter **y** to confirm that you want to continue to install the package.

```

Installing JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster
Edition) as <JNPRsbr>

```

```

## Executing preinstall script.
## Installing part 1 of 1.
.
.
.
[ verifying class <none> ]
## Executing postinstall script.
Newly installed server directory will be backed up as:
/opt/JNPRsbr/radius/install/backups/2009:03:31-00:34:06

```

Installation of <JNPRsbr> was successful.

## Configuring the SBR Carrier Software on the New Management Node

Before starting this procedure, review [“Before You Install Software” on page 39](#). In particular, review requirements for [“Setting Up External Database Connectivity \(Optional\)” on page 53](#) and [“Installing the SIGTRAN Interface \(Optional\)” on page 55](#), because steps in this procedure require the server to be preconfigured for these capabilities.

To configure the software on the new management node:

1. As root, navigate to the **radius/install** subdirectory of the directory where you installed the Steel-Belted Radius Carrier package in [“Installing the SBR Carrier Software on the New Management Node Host Machine” on page 245](#).

Example: `cd /opt/JNPRsbr/radius/install`

2. Run the **configure** script.

Execute:

`./configure`

3. Review and accept the Steel-Belted Radius Carrier license agreement.

Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

4. From the menu of configuration tasks, enter **3** to specify **Configure Cluster Node**.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-8.carrier.spgma.juniper.net
is not configured and processes are down, needs to be configured
-----
```

```
1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.
```

```
2.  Generate Cluster Definition
    Creates new or updates existing cluster definitions.
    Modifies the shared directory but does not modify this node.
```

### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q): 3

## 5. Specify the name of the cluster.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-8.carrier.spgma.juniper.net
is not configured and processes are down, needs to be configured
-----
```

Configuring Cluster Node...

```
Enter SBR cluster name [sbrha]:cambridge
```

```
Reading shared configuration from /opt/JNPRshare/install/cambridge
```

```
Generating configuration files
```

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files: a
```

6. Enter **a** to accept the modified cluster definition files and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

7. The configure script prompts you with a warning whether or not to apply the cluster definition to this node. Enter **y** to proceed.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y
```

```
Cleaning directories
/opt/JNPRhadm
```

```
Applying configuration
```

Initializing Session State Register, please wait a few minutes...

Configuration complete

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,1m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-8.carrier.spgma.juniper.net(s)
is configured and processes are down, may be reconfigured if desired
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

READY: last operation succeeded, node configured.  
Enter the number of the desired configuration task or quit (4,q):

**NOTE:** Expect package is required to be installed on RHEL 7.3 or later (see [“Linux” on page 47](#) for supported RHEL versions). SBR Carrier has been tested only with the Expect-5.45 version. Expect package is installed on Solaris 11.3.36.10.0 or later version by default (see [“Solaris” on page 47](#) for supported Solaris versions).

8. Enter **q** to quit.

### Configuring Each Existing Node in the Cluster with the New Cluster Definition Files

At this point in the process, all nodes in the cluster have the new cluster definition files loaded. However, only the new management node has been configured with the new files. The existing nodes are still running with the old cluster definition files.

In this procedure, you log in to each existing node, stop the processes, run the **configure** script, and restart the processes. You must complete these steps on every existing node in the cluster. This example starts with the sm1 node.



**CAUTION:** In this procedure, you need to stop and restart each existing node one by one to apply the new cluster definition to each of the original cluster nodes. Do not operate on multiple nodes at the same time because that creates multiple faults that can stop the entire cluster.

Always review the recommended start and stop order and processes and plan out the order in which to perform the equivalent steps in your cluster. See [“When and How to Restart Session State Register Nodes, Hosts, and Clusters” on page 196.](#)

1. Log in to the first existing node (in this example, sm1) as root.
2. Navigate to the **radius** subdirectory of the directory in which the **JNPRsbr** package was installed (**/opt/JNPRsbr** by default).  
Example: **cd /opt/JNPRsbr/radius**
3. Stop the RADIUS processes on the node you are configuring (required on each (s) and (sm) node), and execute:
  - a. **./sbrd stop radius**
  - b. **./sbrd status**
4. Stop the **ssr** processes on the node you are configuring (required on each (sm), (m) and (d) node):  
**./sbrd stop ssr**
5. Check the status of the node:  
**./sbrd status**

```
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
id=10   @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11   @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1 (not connected, accepting connect from 172.28.84.36)
id=2   @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)
```

```
[mysqlld(API)] 4 node(s)
id=6 (not connected, accepting connect from 172.28.84.36)
id=7 @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=58 @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=59 (not connected, accepting connect from 172.28.84.36)
```

```
hadm@wrx07:~>
```

6. Verify that the node you are about to configure is not connected. In this example, the node ID for sm1 indicates: **id=1 (not connected, accepting connect from 172.28.84.36)**, indicating the sm1 node is stopped.
7. Run the **configure** script:

Execute:

**./configure**

```
Configuring SBR Software
```

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is configured and processes are down, may be reconfigured if desired
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.



### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (4,q):

8. From the menu of configuration tasks, enter **3** to specify **Configure Cluster Node**.

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is configured and processes are down, may be reconfigured if desired
-----
```

Configuring Cluster Node...

```
Enter SBR cluster name [cambridge]:
```

9. Press Enter to accept the existing cluster name (in the example: cambridge).

```
Create (c) new or update (u) existing node configuration? [u]:
```

10. Enter **u** to update the existing node configuration.

```
Reading shared configuration from /opt/JNPRshare/install/cambridge
```

```
Generating configuration files
```

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files:
```

11. Enter **a** to accept the new cluster definition files.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y
```

The configure script prompts you with a warning about whether or not to apply the new cluster definition to the node.

12. Enter **y** to continue.

```
Applying configuration
```

```

-----
SBR 8.60.50006 cluster cambridge{0s,2sm,1m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is configured and processes are down, may be reconfigured if desired
-----

```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

READY: last operation succeeded, node configured.  
Enter the number of the desired configuration task or quit (4,q):

13. Enter **q** to quit.

```
root@sbrha-4:/opt/JNPRsbr/radius/install>
```

14. Notice that the first line in the applied configuration is: **SBR 8.60.50006 cluster cambridge{0s,2sm,1m,2d}**, indicating that the new configuration has been applied to the sm1 node.

15. Restart the SSR process on the newly configured sm1 node:

- a. Execute:  
**./sbrd start ssr**
- b. Execute:  
**./sbrd status**
- c. Examine each line and ensure the SSR process is running without error.

16. Restart the RADIUS process on sm1:

- a. Execute:  
**./sbrd start radius**
- b. Execute:  
**./sbrd status**
- c. Examine each line and ensure the RADIUS process is running without error.

17. Repeat steps 1 through 16 on the sm2 node and then on each data node, one at a time. For the data nodes, you do not need to stop and restart the RADIUS process because data nodes only run the SSR process.

Do not operate on multiple nodes at once. Doing so creates multiple faults that can stop the entire cluster. Complete the procedure on each node one at a time until the node is operating without error.

## Starting the New Management Node

At this point in the process, the original nodes in the cluster (sm1, sm2, d1, and d2) are up and running with the new cluster definition files. The new management node is configured with the proper configuration, but is not yet running in the cluster. The following procedure starts the new management node in the cluster.

1. Start the `ssr` process on the new management node:
  - a. Log in as root to the management (m) node.
  - b. Change directories to `/opt/JNPRsbr/radius/`.
  - c. Execute:  
`./sbrd start ssr`
  - d. Execute:  
`./sbrd status`
  - e. Examine each line of the final cluster configuration and ensure it is running without error.
2. Now that the new management node is started and running in the cluster, configure it using Web GUI. See [“Basic SBR Carrier Node Configuration” on page 138](#). For complete details, see the *SBR Carrier Administration and Configuration Guide*.

## Running CreateDB.sh on the New Management Node

At this point, the new management node is up and running in the cluster. You run the **CreateDB.sh** script to create certain files that are required to run the administrative shell scripts used to administer the session database in the cluster.

Run the **CreateDB.sh** script on the new management node.

1. Log in as **hadm**.
2. Navigate to the hadm user's home directory, `/opt/JNPRhadm/` by default.
3. Execute:  
**CreateDB.sh**

The addition of the new management node is complete.

## Adding a Data Expansion Kit to an Existing Cluster

Adding the two new data nodes in the Data Expansion Kit to an existing cluster requires deleting and re-creating the session database for the cluster after all the data nodes are up and running.

Because the process of updating the existing cluster topology and re-creating the session database may result in a longer downtime than desired, there are two approaches you can take to minimize the downtime:

- Use a transition server (temporary cluster)—One of the sm nodes in the existing cluster is borrowed from the cluster and converted to a transition server operating a temporary cluster. All traffic is routed to the transition server while the existing cluster is updated to include the new data nodes. After the updated cluster is up and running with the new data nodes, traffic is switched back to it, and the transition server is unconfigured from being a temporary cluster, reconfigured as a sm node, and re-incorporated into the updated cluster. We use this method in this example procedure. See [“Using a Transition Server When Adding Data Nodes to an Existing Cluster” on page 261](#).
- Non-transition server—This approach results in longer downtime, the cluster database is destroyed, and entire cluster is updated and reconfigured with the new topology before re-creating the old database. See [“Non-Transition Server Method—Terminating Connections” on page 307](#).

**NOTE:** Although both of these methods minimize downtime as much as possible, they both require the cluster to be reinitialized, which necessitates destroying and re-creating the session database. The difference between the two approaches is that using a transition server allows SBR Carrier traffic to be processed while the remaining cluster is updated. This is not possible with the non-transition server. See [“Non-Transition Server Method—Terminating Connections” on page 307](#).

## Requirements for Selecting a Transition Server in Your Environment

Use the following selection criteria to select a temporary server:

- The server must meet all the Release 8.6.0 hardware and software requirements listed in [“Before You Install Software” on page 39](#).
- If the server is part of an existing cluster:
  - We recommend using the most powerful (the most RAM and greatest number of processors) available because it will be processing a heavier-than-normal load during the transition.
  - We recommend using a SBR or management node, rather than a data node both to reduce front end processing on the existing cluster and to maintain data redundancy.
- If you intend the server to be the transition server and then reconfigure it to be part of the updated cluster when it is reconfigured, it must be a combined SBR/management node host.
- If you use Centralized Configuration Management to replicate SBR Carrier node configurations among a group of like nodes, the transition server cannot take on the role of primary CCM server in the updated cluster because it will not be the first SBR node to be configured.

## Using a Transition Server When Adding Data Nodes to an Existing Cluster

In general, to use a transition server to add a Data Expansion Kit to an existing cluster:

1. Create the transition server and switch all traffic to it.  
See [“Creating the Transition Server” on page 263](#).
2. Create the updated cluster definition files that include the two new data nodes.  
See [“Creating the Updated Cluster Definition Files” on page 271](#).
3. Install the SBR Carrier software on the host machines for the new data nodes.  
See [“Installing the SBR Carrier Software on the Two New Data Node Host Machines” on page 278](#).
4. Distribute the new cluster definition files to the existing cluster nodes and the new data nodes.  
See [“Distributing the Updated Cluster Definition Files to the Existing Nodes” on page 281](#).
5. Destroy the session database on the existing cluster.  
See [“Destroying the Session Database on the Original Cluster” on page 282](#).
6. Configure each node in the expanded cluster with the updated cluster definition files.  
See [“Configuring the Nodes in the Expanded Cluster with the Updated Cluster Definition Files” on page 284](#).
7. Create the session database and IP pools for the expanded cluster.  
See [“Creating the Session Database and IP Pools on the Expanded Cluster” on page 293](#).
8. Switch the traffic back to the updated, expanded cluster.  
See [“Removing the Transition Server from Service” on page 297](#).
9. Unconfigure the transition server, rebuild it, and reincorporate it into the expanded cluster.  
See [“Unconfiguring and Rebuilding the Transition Server” on page 297](#).

## Existing Cluster Configuration for This Example Procedure

The following procedure adds one Data Expansion Kit to an existing cluster.

The following designations are used throughout the examples in this section:

sm = Hardware has SBR node and Management node.

s = Hardware has only SBR node.

m = Hardware has only Management node.

d = Hardware has Data node.

2sm, 2d = Two SBR/Management nodes and 2 Data nodes.

2S, 2SM, 2D = Two SBR nodes and 2 SBR/Management nodes, 2 Data nodes.

Display the existing cluster:

#### Cluster Configuration

-----

[ndbd(NDB)] 2 node(s)

id=10 @172.28.84.163 (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)

id=11 @172.28.84.113 (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)

[ndb\_mgmd(MGM)] 2 node(s)

id=1 @172.28.84.36 (mysql-5.7.25 ndb-7.6.9)

id=2 @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)

[mysqld(API)] 5 node(s)

id=6 @172.28.84.36 (mysql-5.7.25 ndb-7.6.9)

id=7 @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)

id=58 @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)

id=59 @172.28.84.36 (mysql-5.7.25 ndb-7.6.9)

The existing cluster includes two full SBR Carrier licenses and a license for the Starter Kit resulting in a configuration that includes two sm nodes and two d nodes. For the purposes of this procedure, the existing two sm nodes are identified as sm1 and sm2 as follows:

id=1 @172.28.84.36 = sm1

id=2 @172.28.84.166 = sm2

In this example, we *borrow* the sm2 node, and convert it to a transition server operating as a temporary cluster.



## Creating the Transition Server

In this example, we *borrow* the sm2 node and convert it to a transition server operating as a temporary cluster. To set up the transition server to temporarily take the place of the existing cluster, you need to prepare the server, install software, and configure the database.

The SBRC Temporary Cluster, also termed *the transition server*, is an exceptional node in the sense that it executes all processes on one machine. The transition server is assigned the node type= smdt. The **CreateDB.sh** and configuration of pool(s) must be done manually on a transition server, just like in a cluster.

### Stopping the Processes on the Target Transition Server

1. Log in to the server that you are reconfiguring to act as the transition server (in this example sm2) as root.
2. Navigate to the **radius** subdirectory of the directory in which the **JNPRsbr** package was installed (**/opt/JNPRsbr** by default).

Example: **cd /opt/JNPRsbr/radius**

3. Stop the RADIUS and SSR processes on the node.

Execute:

**./sbrd stop radius**

**./sbrd stop ssr**

4. Check the status of the node and confirm it is not connected to the cluster:

**./sbrd status**

```
Cluster Configuration
```

```
-----
```

```
[ndbd(NDB)]      2 node(s)
id=10   @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11   @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1    @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2 (not connected, accepting connect from 172.28.84.166)
```

```
[mysqld(API)]   4 node(s)
id=6 (not connected, accepting connect from 172.28.84.36)
id=7    @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)
```

```
id=58   @172.28.84.166   (mysql-5.7.25 ndb-7.6.9)
id=59 (not connected, accepting connect from 172.28.84.36)
```

```
hadm@wrx07:~>
```

In this example, notice that the sm2 node is not connected as indicated by

```
id=2 (not connected, accepting connect from 172.28.84.166).
```

### ***Configuring the Software on the Transition Server as a Temporary Cluster***

Now that the processes are stopped on the machine we are reconfiguring as the transition server, we need to reconfigure it as a temporary cluster. At this point, you are still logged in to the target machine as root (in this case the original sm2 node).

1. Execute the **configure** script to reconfigure the machine as a temporary cluster:

Execute:

**./configure**

Example:

```
root@wrx07:/opt/JNPRsbr/radius/install> ./configure
```

```
Configuring SBR Software
```

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node wrx07(sm)
is not configured and processes are down, needs to be configured
-----
```

1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on any node.

Must be invoked on each and every node of the cluster.

4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q):

2. From the menu of configuration tasks, enter **5** to specify **Create Temporary Cluster**.

```
Creating Temporary Cluster...
```

```
Enter SBR cluster name [wrx07]:
```

3. Enter the exact name of the existing cluster. In this example: **cambridge**.

```
In order to avoid service outages when performing certain major
cluster maintenance tasks, you are allowed to reuse each of
your licenses in order to create a temporary cluster that
consists of 1 SBR node, 1 management node, and 1 data node all
installed on the same machine. Note that this is not a true
cluster since it is vulnerable to single points of failure.
```

4. Enter the SSR Starter Kit license number, the license number for one SBR node, and, if you are using one of the optional SBR Carrier modules, the license number for it.

While migrating to the updated cluster, you can use the same licenses for the transition server as for the updated cluster.

```
Enter Starter Kit license: XXXX XXXX XXXX XXXX
Enter SBR licenses meant only for this particular SBR node.
Enter one license per line and an empty line when finished.
Enter SBR full license: XXXX XXXX XXXX XXXX
Enter SBR feature license:
```

5. Enter passwords for two internal accounts. The password input is not echoed to the screen; the fields appear to be blank.

```
All cluster nodes will share the same Session State Register (SSR).
Setting password for SSR admin account hadmsql
Password:
Again:
Setting password for SSR software account hadmsbr
Password:
Again:
```

```
Generating configuration files
```

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files:
```

The system generates the required configuration files and prompts you to view, accept, or reject them.

6. Enter **a** to accept them and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n):
```

You are prompted with a warning whether or not to apply the changes.

7. Enter **y** to continue.
8. For the remainder of the prompts, simply press Enter to configure the transition server with the existing configuration.

```
Cleaning directories
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRmysqld
/opt/JNPRndb_mgmd
/opt/JNPRndbd
```

```
Applying configuration
```

```
Initializing Session State Register, please wait a few minutes...
```

```
Configuring for use with generic database
Do you want to configure Java Runtime Environment for JDBC Feature [n] :
Do you want to enable "Radius WatchDog" Process? [n]:
Do you want to enable LCI? [n]:
Do you want to configure for use with Oracle? [n]:
Removing oracle references from startup script
Do you want to configure for use with SIGTRAN? [n]:
Removing SIGTRAN references from startup script
Do you want to configure SNMP? [n]:
Configuring Admin GUI Webserver
Compatible Java version found :
Do you want to install custom SSL certificate for Admin WebServer? [n]:
Enable (e), disable (d), or preserve (p) autoboot scripts [e]:
```

```
The SBR Admin Web GUI can be launched using the following URL:
https://<servername>:2909
```

```
Configuration complete
```

```
-----
SBR 8.60.50006 temporary cluster cambridge
on SunOS 5.10 Generic_141444-09 node wrx07(smdt)
is configured and processes are down, may be reconfigured if desired
-----
```

#### 1. Unconfigure Cluster Node

```
Not used when merely updating existing cluster definitions.
```

## 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.

## 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

## 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

## 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

## 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

READY: last operation succeeded, created temporary cluster.  
Enter the number of the desired configuration task or quit (4,q):

9. Enter **q** to quit.

10. Notice the server configuration in the line:

```
SBR 8.60.50006 temporary cluster cambridge
on SunOS 5.10 Generic_141444-09 node wrx07(smdt)
```

(smdt) indicates the machine is configured as an s,m,d temporary cluster.

### **Configuring and Starting the Transition Server**

Now that the software is configured, you need to create the session database and the IP pools and ranges on the transition server. All cluster traffic will ultimately be switched to this single transition server temporarily, while you take the other nodes in the existing cluster down and upgrade and reconfigure them. So, you need to configure the temporary transition server to match the existing cluster configuration.

1. Navigate to the **radius** subdirectory of the directory in which the **JNPRsbr** package was installed (/opt/JNPRsbr by default) and start the **SSR** process on the transition server.

Example: **cd /opt/JNPRsbr/radius**

2. As root, execute:

**./sbrd start ssr**

Status messages are displayed as the programs start:

```
Starting ssr management processes
Starting ssr auxiliary processes
Starting ssr data processes
```

3. Verify the process started without error:

As root, execute:

**./sbrd status**

4. Create the session database.

If you need to customize the sessions database to match your existing cluster session database, see [“Customizing the SSR Database Current Sessions Table” on page 144](#). Any customization must be done prior to running the **CreateDB.sh** script.

- a. Log in as **hadm**.
- b. Navigate to the hadm user's home directory, **/opt/JNPRhadm/** by default.

Execute:

**./CreateDB.sh**

5. As hadm, set up IP address pools and ranges using the SSR Administration Scripts. The IP address range should be separate from the in-use pools on the existing and upgraded cluster to avoid overlaps. If the old and transitional pools overlap, then during the transition the two clusters may give the same IP



address to two different users. See the section *Session State Register Administration* in the *SBR Carrier Administration and Configuration Guide* for more information.

6. Start the RADIUS process:

As root, execute:

**sbrd start radius**

Status messages are displayed as the programs start:

```
Starting radius server processes
RADIUS: Process ID of daemon is 13224
RADIUS: Starting DCF system
RADIUS: Configuration checksum: 2D D6 38 1D
RADIUS started
.
.
.
RADIUS: DCF system started
```

7. Verify the process started without error:

As root, execute:

**./sbrd status**

8. Finish configuring the transition server using Web GUI. Follow the steps outlined in [“Basic SBR Carrier Node Configuration” on page 138](#). For complete details, see the *SBR Carrier Administration and Configuration Guide*.

### **Switching Traffic to the Transition Server**

After the transition server is set up and tested, and a working database created, reconfigure the site's routers to gradually direct traffic to the transition server instead of to the existing cluster's SBR servers.

### **Creating the Updated Cluster Definition Files**

The next phase of the process is to create the new cluster definition files to include the two new data nodes from the Data Expansion Kit. At this point in the process the existing cluster configuration shows the sm2 node processes are not running and not connected, as indicated by **id=2 (not connected, accepting connect from 172.28.84.166)**:

```
Cluster Configuration
-----
[nbdb(NDB) ]      2 node(s)
```

```
id=10   @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11   @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1     @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2 (not connected, accepting connect from 172.28.84.166)
```

```
[mysqld(API)] 5 node(s)
id=6     @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=7 (not connected, accepting connect from 172.28.84.166)
id=58 (not connected, accepting connect from 172.28.84.166)
id=59    @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
```

Start by creating the updated cluster definition files on the sm1 node:

1. As root, on the sm1 node, navigate to the **radius/install** subdirectory of the directory in which the **JNPRsbr** package was installed (**/opt/JNPRsbr** by default).

Example: **cd /opt/JNPRsbr/radius/install**

2. Run the **configure** script:

Execute:

**./configure**

Example:

**root@sbrha-4:/opt/JNPRsbr/radius/install> ./configure**

Configuring SBR Software

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
```

is CONFIGURED and processes are UP, may be stopped if reconfigured

---

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

```
Enter the number of the desired configuration task or quit (4,q):
```

- From the menu of configuration tasks, enter **2** to specify **Generate Cluster Definition**.

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is CONFIGURED and processes are UP, may be stopped if reconfigured
-----
```

```
Generating Cluster Definition...
```

```
Enter SBR cluster name [cambridge]:
```

You are prompted to enter the name of the cluster.

- Press Enter to use the current cluster name.

You are prompted either to create a new cluster or update an existing cluster definition.

```
Create (c) new or update (u) existing cluster definition? [u]:
```

- Enter **u** to update the existing cluster. definition.

```
The SBR Cluster Starter Kit license allows you to create a minimal
cluster of 2 SBR nodes, 2 management nodes, and 2 data nodes.  When
each node is installed on a separate machine the cluster topology is
denoted as {2s,2m,2d}.  When SBR nodes are paired with management
nodes on the same machines the cluster topology is denoted as {2sm,2d}.
```

```
An optional SBR Cluster Management Expansion Kit allows you to add a
third management node for {2sm,1m,2d} and an optional Data Expansion
```

Kit allows you to add 2 more data nodes for {2sm,1m,4d} clusters.  
 Additional SBR licenses allow you to add up to 18 more SBR nodes to  
 obtain a maximal cluster {18s,2sm,1m,4d} and/or enable extra features.

While it is not difficult to add management and/or SBR nodes to an  
 existing cluster, adding data nodes is more difficult and may require  
 you to shutdown the entire cluster as opposed to a rolling restart.

Another license is required if you wish to add a third management node.  
 Adding a third management node will require a rolling restart later.  
 Enter Management Expansion Kit license, if any:

6. Because we are not adding a Management Expansion Kit, press Enter to skip adding the license.

Another license is required if you wish to add more data nodes.  
 Adding data nodes may require you to shutdown the entire cluster.  
 Enter Data Expansion Kit license, if any: **XXXX XXXX XXXX XXXX**

7. Enter the license number for the Data Expansion Kit and press Enter.

This cluster presently contains 2 of 20 possible SBR nodes.  
 Adding more SBR nodes will require a rolling restart later.  
 Enter number of SBR nodes to be added [0]:

8. When prompted to enter the number of SBR nodes, press Enter to keep the existing configuration.

Updating cluster cambridge{0s,2sm,0m,4d}  
 will require 2 new machines. Do you wish to continue? [y]:

9. Notice the updated cluster configuration includes four data nodes as indicated by: **Updating cluster cambridge{0s,2sm,0m,4d}**.

Enter y to continue.

10. When prompted, enter the node names and IP addresses for the two new data nodes.

Press Enter when prompted to Enter node type (d) [d]: and when prompted to **Enter DATA node ID**.

Information will now be gathered for each new machine to be added.  
 You will have a chance to review all information at least once  
 before any machines are modified.

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,4d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
Partial configuration at present is {0s,2sm,0m,2d} of {0s,2sm,0m,4d}
-----
IMPORTANT: node names must be entered as reported by 'uname -n'.
```

```
Enter node name [cambridge-6]: sbrha-8.carrier.spgma.juniper.net
Enter node type (d) [d]:
Enter DATA node ID (10-29) [12]:
Enter DATA node IP address by which it is known to management nodes.
Enter DATA node IP address: 172.28.84.178
-----
```

```
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,4d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
Partial configuration at present is {0s,2sm,0m,3d} of {0s,2sm,0m,4d}
-----
IMPORTANT: node names must be entered as reported by 'uname -n'.
Enter node name [cambridge-7]: sbrha-2.spgma.juniper.net
Enter node type (d) [d]:
Enter DATA node ID (10-29) [13]:
Enter DATA node IP address by which it is known to management nodes.
Enter DATA node IP address: 172.28.84.104
```

The system generates the updated cluster definition files.

11. Verify the proper configuration by examining the line: **Generated configuration is {0s,2sm,0m,4d} of {0s,2sm,0m,4d}** showing the four data nodes.

When prompted to, enter **a** to accept the updated configuration.

```

-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,4d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
Generated configuration is {0s,2sm,0m,4d} of {0s,2sm,0m,4d}
-----

```

Generating configuration files

```

Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files: a
Writing shared configuration to /opt/JNPRshare/install/cambridge

```

```

-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is CONFIGURED and processes are UP, may be stopped if reconfigured
-----

```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

READY: last operation succeeded, generated cluster definition.  
Enter the number of the desired configuration task or quit (4,q):

12. When the main configuration menu is displayed, enter **q** to quit.

```
root@sbrha-4:/opt/JNPRsbr/radius/install>
```

## Installing the SBR Carrier Software on the Two New Data Node Host Machines

At this point in the process, the updated cluster definition files have been generated and reside on the sm1 node only. Next you need to install the SBR Carrier software on each of the machines that you want to host the two new data nodes. After the SBR Carrier software is installed on these machines, you distribute the updated cluster definition files to all the other nodes in the original cluster.



This procedure describes how to unpack and install the SBR Carrier software on the host machine for the new data nodes.

1. Log in to the machine as root.
2. Copy the Steel-Belted Radius Carrier installation files from their download location to the machine. Make sure to copy them to a local or remote hard disk partition that is readable by root.

This example copies the files from a download directory to the **/tmp/sbr** directory.

Execute:

```
mkdir -p /opt/tmp  
cp -pR /tmp/sbr/solaris/* /opt/tmp/
```

3. Extract the SBR Carrier installation package.

For 64-bit Solaris, execute:

```
cd /tmp/sbr  
ls -ltr
```

```
total 216240  
-rw-r--r--  1 root    root      110712276 Aug 25 09:44  
sbr-cl-8.6.0.R-1.sparcv9.tgz
```

Execute:

```
gunzip sbr-cl-8.6.0.R-1.sparcv9.tgz  
tar xf sbr-cl-8.6.0.R-1.sparcv9.tar
```

4. Verify that the extraction worked and confirm the name of the package file.

For 64-bit Solaris, execute:

```
ls -ltr
```

```
total 216256  
drwxr-xr-x  4 Xtreece  other      370 Aug 24 17:01 JNPRsbr.pkg  
-rw-r--r--  1 root    root      110712276 Aug 25 09:44  
sbr-cl-8.6.0.R-1.sparcv9.tar
```

5. Install the package.

Execute:

```
pkgadd -d /tmp/sbr
```

```
The following packages are available:  
1  JNPRsbr.pkg          JNPRsbr - Juniper Networks Steel-Belted Radius
```

```
(Carrier Cluster Edition)          (sparc) 8.60.50006
```

```
Select package(s) you wish to process (or 'all' to process all packages).
(default: all) [?,??,q]:
```

6. Type **all** and press Enter.

The script resumes.

```
Processing package instance <JNPRsbr.pkg> from </tmp>
```

7. Confirm the installation directory.

Depending on the system configuration, you are prompted whether to create the **/opt/JNPRsbr** directory if it does not exist, to over-write an already extracted package, or any of several other questions.

```
The selected base directory </opt/JNPRsbr> must exist before installation is
attempted.
```

```
Do you want this directory created now [y,n,?,q]
```

8. Answer the question appropriately (or change the extraction path if necessary) so that the script can proceed.

To accept the default directory as a target, enter **y**.

The script resumes.

```
Using </opt/JNPRsbr> as the package base directory.
#Processing package information.
#Processing system information.
    48 package pathnames are already properly installed.
#Verifying disk space requirements.
#Checking for conflicts with packages already installed.
#Checking for setuid/setgid programs.
```

```
This package contains scripts which will be executed with super-user
permission during the process of installing this package.
```

```
Do you want to continue with the installation of <JNPRsbr> [y,n,?]
```

9. Enter **y** to confirm that you want to continue to install the package.

```
Installing JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster
Edition) as <JNPRsbr>
```

```
## Executing preinstall script.
## Installing part 1 of 1.
.
.
.
[ verifying class <none> ]
## Executing postinstall script.
Newly installed server directory will be backed up as:
/opt/JNPRsbr/radius/install/backups/2009:03:31-00:34:06
```

```
Installation of <JNPRsbr> was successful.
```

10. Repeat this process on the second new data node.

## Distributing the Updated Cluster Definition Files to the Existing Nodes

Now that the two machines hosting the new data nodes have the SBR Carrier software installed, you can distribute the updated cluster definition files to the new nodes and the other nodes in the original cluster.

On both the existing nodes and the new data nodes in the original cluster, create a copy of the new cluster definition files. This process does not invoke the updated cluster definition files, but makes them available to the **configure** script later in the workflow.

To distribute the new cluster definition files:

1. Log in to each node (existing and new) as **hadm**.
2. Change directories to the install directory.

(On new nodes, the entire path may not exist because the **<cluster name>** portion of the path was not created when you prepared the new machine, so you may need to create it.) See [“Creating Share Directories” on page 52](#).

Execute:

```
cd /opt/JNPRshare/install/ <cluster_name>
```

For example:

```
cd /opt/JNPRshare/install/cambridge
```

3. Use FTP binary mode to connect to the node host (in this example, sm1) where you created the new cluster definition files.
4. Execute the **get** command to transfer the **configure. <cluster name> .tar** file.

For example:

```
bin
```

```
get /opt/JNPRsbr/radius/install/configure.cambridge.tar
```

5. In a terminal window, extract the new cluster definition files from the archive.

Execute:

```
tar xvf configure. <cluster_name> .tar
```

Output similar to this example is displayed:

```
$ tar xvf configure.MyCluster.tar
x dbcluster.rc, 1925 bytes, 4 tape blocks
x config.ini, 2435 bytes, 5 tape blocks
x my.cnf, 1017 bytes, 2 tape blocks
x dbclusterndb.gen, 33474 bytes, 66 tape blocks
x dbcluster.dat, 921 bytes, 2 tape blocks
```

6. Repeat these steps until every node in the cluster has a copy of the new cluster definition files.

## Destroying the Session Database on the Original Cluster

You now log in to the sm1 node, destroy the session database from the original cluster, and stop the original cluster.

1. Log in to sm1 as **hadm**.
2. Navigate to the hadm user's home directory, **/opt/JNPRhadm** by default.

3. Execute:

**/DestroyDB.sh**

```
hadm@sbrha-4:~> ./DestroyDB.sh
SBRs must be offline; OK? <yes|no>  yes
This will destroy the "SteelBeltedRadius" database; OK? <yes|no>  yes
Really? <yes|no>  yes
```

4. Each time you are prompted as to whether you really want to destroy the database, enter **yes**.

The system responds with:

```
Database "SteelBeltedRadius" destroyed.
```

5. Stop the original cluster by executing:

**/sbrd stop cluster**

```
hadm@sbrha-4:~> su
Password:
# bash
root@sbrha-4:~>
root@sbrha-4:~> /etc/init.d/sbrd stop cluster
WARNING: This function is capable of stopping multiple nodes.
Do not use this function if you intend to stop only one node.
Do you intend to stop the entire cluster? (y,n): y
Are you sure? (y,n): y
Really? (y,n): y
```

6. Each time you are prompted as to whether you really want to stop the entire cluster, enter **y**.

The software stops the RADIUS processes first and then the SSR processes.

```
Stopping radius server processes
waiting for radius
10 seconds elapsed, still waiting
radius stopped
Stopping ssr auxiliary processes
Stopping ssr management processes
Connected to Management Server at: 172.28.84.36:5235
Shutdown of NDB Cluster node(s) failed.
* 1006: Illegal reply from server
```

```
*
root@sbrha-4:~>
```

7. On each remaining node of the original cluster, execute **/sbrd stop cluster** and verify that the processes are stopped.

Perform this step on the remaining nodes in this order: s nodes, sm nodes, m nodes, d nodes.

- a. Log in to each remaining node in the existing cluster as root.
- b. Navigate to the **radius** subdirectory of the directory in which the **JNPRsbr** package was installed (**/opt/JNPRsbr** by default).

Example: **cd /opt/JNPRsbr/radius**

- c. Execute:  
**/sbrd stop cluster**
- d. Execute:  
**/sbrd status**
- e. Examine each line to ensure it says **not connected**.

## Configuring the Nodes in the Expanded Cluster with the Updated Cluster Definition Files

To configure the nodes in the expanded cluster with the updated cluster definition files, you run the **configure** script on each node. First you run the script on the two new data nodes, then run it on the original nodes in the cluster (except for the sm2 node, which is still operating as the transition server).

### Configuring the SBR Carrier Software on the New Data Nodes

Configure the software on each new data node:

1. As root, navigate to the directory where you installed the Steel-Belted Radius Carrier package in [“Installing the SBR Carrier Software on the Two New Data Node Host Machines” on page 278](#).

Then, navigate to the **radius/install** subdirectory.

Example: **cd /opt/JNPRsbr/radius/install**

2. Run the **configure** script.

Execute:  
**./configure**

```
# ./configure
Configuring SBR Software
```

3. The End User License Agreement is displayed. Review the Steel-Belted Radius Carrier license agreement.

Press the spacebar to move from one page to the next.

4. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

5. From the menu of configuration tasks, enter **3** to specify **Configure Cluster Node**.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-2.spgma.juniper.net
is not configured and processes are down, needs to be configured
-----
```

```
1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.
```

```
2.  Generate Cluster Definition
    Creates new or updates existing cluster definitions.
    Modifies the shared directory but does not modify this node.
```

```
3.  Configure Cluster Node
    To be preceded by 'Generate Cluster Definition' on one node.
    Must be invoked on each and every node of the cluster.
```

```
4.  Reconfigure RADIUS Server
    Only on SBR nodes, updates the existing SBR configuration.
```

```
5.  Create Temporary Cluster
    Used to approximate a cluster using only this one machine.
    Intended for migration and demonstration purposes only.
```

## 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q): 3

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-2.spgma.juniper.net
is not configured and processes are down, needs to be configured
-----
```

Configuring Cluster Node...

Enter SBR cluster name [sbrha]: cambridge

## 6. Enter the exact name of the cluster and press Enter.

Reading shared configuration from /opt/JNPRshare/install/cambridge

Generating configuration files



```

Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
View (v), accept (a), or reject (r) configuration files:

```

7. Enter **a** to accept the configuration.

```

WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n):

```

8. Enter **y** to continue.

```

Cleaning directories
/opt/JNPRhadm

```

```

Applying configuration

```

```

Initializing Session State Register, please wait a few minutes...

```

```

-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,4d}
on SunOS 5.10 Generic_141444-09 node sbrha-2.spgma.juniper.net(d)
is configured and processes are down, may be reconfigured if desired
-----

```

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.

## 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.

## 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

## 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

## 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

## 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

READY: last operation succeeded, node configured.  
Enter the number of the desired configuration task or quit (2,q):

9. Enter **q** to quit.

10. Notice the line: **node sbrha-2.spgma.juniper.net(d) is configured and processes are down, may be reconfigured if desired** indicating the node name you assigned and that the node was configured without error. The processes remain down for now.

11. Log in to the next new data node and repeat this procedure.

### *Running the Configure Script on Each Node from the Original Cluster*

At this point in the process, all nodes in the cluster have the new cluster definition files loaded. However, only the new data nodes have been configured with the new files.

In this step, you run the configure script on each node from the original cluster. This includes the sm1, d1, and d2 nodes. Running this script applies the updated cluster definition files to the nodes.

You do not run the script on the sm2 node, which is still operating as the transition server (temporary cluster).

1. Log in to the first existing node (in this example, sm1) as root.
2. Navigate to the **radius** subdirectory of the directory in which the JNPRsbr package was installed.  
Example: `cd /opt/JNPRsbr/radius/`
3. Check the status of the node by executing:  
`./sbrd status`
4. Examine the line for the node you are about to configure, and verify that it is not connected. In this example, the node ID for sm1 indicates: **id=1 (not connected, accepting connect from 172.28.84.36)**, indicating the sm1 node is stopped.
5. Navigate to the **radius/install** subdirectory of the directory where the JNPRsbr package was installed.  
Example: `cd /opt/JNPRsbr/radius/install`
6. Run the **configure** script to apply the updated cluster definition files:

Execute:

`./configure`

```
root@sbrha-4:/opt/JNPRsbr/radius/install> ./configure
Configuring SBR Software
```

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is configured and processes are down, may be reconfigured if desired
-----
```

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.

## 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.

## 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

## 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

## 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

## 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (4,q):

7. Enter **3** to specify **Configure Cluster Node** and press Enter.

```
-----  
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
```

```
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is configured and processes are down, may be reconfigured if desired
-----
```

```
Configuring Cluster Node...
```

```
Enter SBR cluster name [cambridge]:
```

8. Press Enter to accept the cluster name and continue.

You are prompted either to create a new or update an existing node configuration.

```
Create (c) new or update (u) existing node configuration? [u]:
```

9. Enter **u** to update the node with the updated cluster definition files.

```
Reading shared configuration from /opt/JNPRshare/install/cambridge
```

```
Generating configuration files
```

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files:
```

10. Enter **a** to accept the updated configuration.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n):
```

## 11. Enter y to continue.

Applying configuration

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,4d}
on SunOS 5.10 Generic_141444-09 node sbrha-4(sm)
is configured and processes are down, may be reconfigured if desired
-----
```

## 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

## 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

## 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

## 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

## 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

```
READY: last operation succeeded, node configured.  
Enter the number of the desired configuration task or quit (4,q):
```

Notice the applied configuration includes the four data nodes as indicated by the line: SBR 8.60.50006 cluster cambridge{0s,2sm,0m,4d}.

12. Enter **q** to quit.

13. Log in to the remaining nodes from the original cluster (d1 and d2) and repeat this procedure.

## Creating the Session Database and IP Pools on the Expanded Cluster

At this point in the process, all nodes in the expanded cluster have been configured with the updated cluster definition files. All of these nodes are currently down. You now create the session database and IP pools and ranges for the expanded cluster. To create the new session database, we recommend that you run the **clean** command on the nodes from the original cluster (in this case, sm1, d1, and d2).

The sm2 node is still operating as the transition server (temporary cluster). Do not disrupt it in any way.

The following procedure describes how to run the **clean** command on sm1, d1, and d2, start the SSR process and create the session database and IP pools.

### *Cleaning the Original Nodes from the Cluster*

Perform the following procedure on sm1, d1, and d2 only:

1. Log in to the first existing node (in this example, sm1) as root.
2. Navigate to the **radius** subdirectory of the directory in which the **JNPRsbr** package was installed.

Example: **cd /opt/JNPRsbr/radius**

3. Execute:  
**./sbrd clean**

```

WARNING: Cleaning the SSR lock on this node may be destructive.
Do not use this function unless you are attempting to start the
entire cluster for the first time, or for recovery purposes.
Clean the SSR lock on this node? (y,n): y
Are you sure? (y,n): y Really? (y,n): y
Cleaning SSR lock

```

4. Repeat this procedure on the d1 and d2 nodes.

### ***Creating the Session Database and IP Pools***

In this procedure, you create the session database and IP address pools for the expanded cluster. For details on performing these tasks, see the section on *Session State Register Administration* in the *SBR Carrier Administration and Configuration Guide*.

First you start the SSR process. The proper order for starting the SSR process is (sm) nodes, (m) nodes, and (d) nodes. We do not have any (m) nodes in this example, so start the SSR process in the following order: sm1, d1, d2, d3, and d4. Start the SSR process on each node in the expanded cluster *one at a time*, starting with the sm1 node and then on each data node. For complete details on the proper order of starting and stopping nodes, see

[“When and How to Restart Session State Register Nodes, Hosts, and Clusters” on page 196.](#)

### ***Starting the SSR Processes on the Nodes in the Expanded Cluster***

1. Log in to the first sm node (in this example, sm1) as root.
2. Navigate to the radius subdirectory of the directory in which the **JNPRsbr** package was installed (/opt/JNPRsbr by default).

Example: `cd /opt/JNPRsbr/radius`

3. Start the SSR process:

`./sbrd start ssr`

4. Before moving on to the next node, verify the SSR process started without error by executing:

`./sbrd status`

5. Examine the status and ensure there are no errors.
6. Repeat this procedure on the d1 and d2 nodes.

When you finish starting the SSR process on sm1, d1, d2, d3, and d4, the cluster configuration is as follows:

```

Cluster Configuration
-----
[ndbd(NDB)]      4 node(s)
id=10    @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)

```



```
id=11    @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
id=12    @172.28.84.178  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 1)
id=13    @172.28.84.104  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 1)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1      @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2 (not connected, accepting connect from 172.28.84.166)
```

```
[mysqld(API)] 4 node(s)
id=6 @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=7 (not connected, accepting connect from 172.28.84.166)
id=58 (not connected, accepting connect from 172.28.84.166)
id=59 (not connected, accepting connect from 172.28.84.36)
```

The lines for node IDs 10, 11, 12, and 13 indicate the SSR processes started without error on the four data nodes.

The line **id=1 @172.28.84.36 (mysql-5.7.25 ndb-7.6.9)** indicates the SSR process started properly on the sm1 node.

Notice that the sm2 node still says it is not connected as indicated by the line: **id=2 (not connected, accepting connect from 172.28.84.166)**. The sm2 node is still operating as the transition server.

### ***Creating the Session Database and IP Address Pools***

Now create the session database and IP pools and ranges on the sm1 node.

1. Log back in to the sm1 node as **hadm**.
2. Navigate to the hadm user's home directory, **/opt/JNPRhadm** by default.
3. Create the session database on the sm1 node.

If you need to customize the sessions database, see

[“Customizing the SSR Database Current Sessions Table” on page 144](#). Any customization must be done before running the **CreateDB.sh** script.

- a. Log in as **hadm**.
- b. Navigate to the hadm user's home directory, **/opt/JNPRhadm/** by default.
- c. Execute:

**./CreateDB.sh**

4. As hadm, add the IP address pools and ranges.

For details on performing these tasks, see the section on *Session State Register Administration* in the *SBR Carrier Administration and Configuration Guide*.

### **Starting the RADIUS Process**

Start the RADIUS process on the sm1 node.

Execute:

**./sbrd start radius**

**./sbrd status**

When you finish starting the SSR process on sm1, d1, d2, d3, and d4, the cluster configuration is as follows:

#### Cluster Configuration

-----

[ndbd(NDB)] 4 node(s)

id=10 @172.28.84.163 (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)

id=11 @172.28.84.113 (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)

id=12 @172.28.84.178 (mysql-5.7.25 ndb-7.6.9, Nodegroup: 1)

id=13 @172.28.84.104 (mysql-5.7.25 ndb-7.6.9, Nodegroup: 1)

[ndb\_mgmd(MGM)] 2 node(s)

id=1 @172.28.84.36 (mysql-5.7.25 ndb-7.6.9)

id=2 (not connected, accepting connect from 172.28.84.166)

[mysqld(API)] 4 node(s)

id=6 @172.28.84.36 (mysql-5.7.25 ndb-7.6.9)

id=7 (not connected, accepting connect from 172.28.84.166)

id=58 (not connected, accepting connect from 172.28.84.166)

id=59 @172.28.84.36 (mysql-5.7.25 ndb-7.6.9)

The RADIUS process for the sm1 node has started properly as indicated by the line:

```
id=59 @172.28.84.36 (mysql-5.7.25 ndb-7.6.9)
```

Notice that the sm2 node is the only node that still is not connected, as indicated by the lines:

```
id=2 (not connected, accepting connect from 172.28.84.166)
id=7 (not connected, accepting connect from 172.28.84.166)
id=58 (not connected, accepting connect from 172.28.84.166)
```

Now that the expanded cluster nodes sm1, d1, d2, d3, and d4 are all started and running without error you can switch traffic back to the expanded cluster.

## Removing the Transition Server from Service

After you bring the expanded cluster online, configure it, and test it, begin transferring live traffic to it and away from the transition server. When all traffic has been shifted to the new expanded cluster and the number of on-going sessions managed by the transition server has reached a suitably low level, take the transition server offline. Some sessions are terminated, but reconnect through the new cluster.

## Unconfiguring and Rebuilding the Transition Server

To free the licenses used by the transition server (in this case, sm2), and clean up installed software, uninstall the SBR Carrier software. See [“Uninstalling Steel-Belted Radius Carrier Software” on page 358](#).

### *Unconfiguring the Transition Server*

1. Log in to the sm2 node as root.
2. Navigate to the **radius** subdirectory of the directory in which the **JNPRsbr** package was installed (/opt/JNPRsbr by default).

Example: `cd /opt/JNPRsbr/radius`

3. Stop the RADIUS processes.

Execute:

```
./sbrd stop radius
```

4. Stop the **SSR** processes:

Execute:

**./sbrd stop ssr**

5. Check the status on the sm2 node to ensure the processes are stopped.

Execute:

**./sbrd status**

6. Navigate to the directory where you installed the SBR Carrier package and then into the **radius/install** subdirectory. (**/opt/JNPRsbr/radius/install**)

7. Run the **unconfigure** script:

Execute:

**./unconfigure**

```
root@wrx07:/opt/JNPRsbr/radius/install> ./unconfigure
Unconfiguring SBR Software
```

```
-----
SBR 8.60.50006 temporary cluster cambridge
on SunOS 5.10 Generic_141444-09 node wrx07(smdt)
is configured and processes are down, may be reconfigured if desired
-----
```

```
Unconfiguring Cluster Node...
```

```
WARNING: You are about to unconfigure this node.
Are you sure that you wish to continue? (y,n):
```

8. At the warning message enter **y** to continue.

```
Cleaning directories
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRmysqld
```

```
/opt/JNPRndb_mgmd
/opt/JNPRndbd
```

```
Locating shared directory...
drwxrwxr-x  2 hadm      hadmg      512 Apr 15 20:19
/opt/JNPRshare/install/cambridge
```

```
WARNING: If you remove the shared directory for this cluster, you
will either have to recover the data from another cluster node or
reconfigure the entire cluster again. This is neither necessary
nor recommended if you are updating an existing configuration.
Remove the shared directory for this cluster? [n]:
```

9. Press Enter indicating you do not want to remove the shared directory.

```
Locating OS user account and home directory...
hadm:x:16663:65536::/opt/JNPRhadm:/bin/bash
hadmg::65536:
drwxrwx---  2 hadm      hadmg      1536 Apr 16 00:03 /opt/JNPRhadm
```

```
WARNING: If you remove the OS user account hadm you will have
to recreate it, the associated OS group account hadmg, and
the associated home directory /opt/JNPRhadm
This is neither necessary nor recommended if you are updating
an existing configuration. Remove the OS user account? [n]:
```

10. Press Enter indicating you do not want to remove the OS user account.

```
Unconfigured
```

```
root@wrx07:/opt/JNPRsbr/radius/install>
```

### Retrieving the Updated Cluster Definition Files from SM1 Node

To distribute the new cluster definition files:

1. Log in to the sm2 node as **hadm**.

2. Change directories to the install directory.

(On new nodes, the entire path may not exist because the **<cluster name>** portion of the path was not created when you prepared the new machine, so you may need to create it.) See [“Creating Share Directories” on page 52](#).

Execute:

```
cd /opt/JNPRshare/install/ <cluster_name>
```

For example:

```
cd /opt/JNPRshare/install/cambridge
```

3. Use FTP binary mode to connect to the node host (in this example, sm1) where you created the new cluster definition files.
4. Execute the **get** command to transfer the **configure. <cluster name> .tar** file.

For example:

```
bin
```

```
get /opt/JNPRsbr/radius/install/configure.cambridge.tar
```

5. In a terminal window, extract the new cluster definition files from the archive.

Execute:

```
tar xvf configure. <cluster_name> .tar
```

Output similar to this example is displayed:

```
$ tar xvf configure.MyCluster.tar
x dbcluster.rc, 1925 bytes, 4 tape blocks
x config.ini, 2435 bytes, 5 tape blocks
x my.cnf, 1017 bytes, 2 tape blocks
x dbclusterndb.gen, 33474 bytes, 66 tape blocks
x dbcluster.dat, 921 bytes, 2 tape blocks
```

### Running the Configure Script on the SM2 Node

1. Log in to the sm2 node as root.
2. Navigate to the **radius/install** subdirectory of the directory in which the JNPRsbr package was installed (**/opt/JNPRsbr** by default).

Example: **cd /opt/JNPRsbr/radius/install**

3. Run the **configure** script to apply the updated cluster definition files:

Execute:

**./configure**

4. Review and accept the Steel-Belted Radius Carrier license agreement.

Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

```
-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node wrx07(sm)
is not configured and processes are down, needs to be configured
-----
```

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.

2. Generate Cluster Definition  
Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.

3. Configure Cluster Node  
To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

4. Reconfigure RADIUS Server  
Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q):

5. From the menu of configuration tasks, enter **3** to specify **Configure Cluster Node**.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node wrx07(sm)
is not configured and processes are down, needs to be configured
-----
```

Configuring Cluster Node...

Enter SBR cluster name []: **cambridge**

6. Specify the exact name of the cluster.

Reading shared configuration from /opt/JNPRshare/install/cambridge



```
Generating configuration files
```

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files:
```

7. Enter **a** to accept the modified configuration files and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

8. The configure script prompts you with a warning whether or not to apply the cluster definition to this node. Enter **y** to proceed.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y
```

```
Cleaning directories
/opt/JNPRhadm
```

```
Applying configuration
```

```
Initializing Session State Register, please wait a few minutes...
```

**NOTE:** Expect package is required to be installed on RHEL 7.3 or later (see [“Linux” on page 47](#) for supported RHEL versions). SBR Carrier has been tested only with the Expect-5.45 version. Expect package is installed on Solaris 11.3.36.10.0 or later version by default (see [“Solaris” on page 47](#) for supported Solaris versions).

## 9. Configure the node.

For information about configuring the node in the following prompts, see [“Configuring the Host Software on the First Server in the Cluster” on page 99](#).

```
Do you want to configure Java Runtime Environment for JDBC Feature [n] :
Please enter backup or radius directory from which to migrate.
Enter n for new configuration, s to search, or q to quit
[n]:
```

```
Enter initial admin user (UNIX account must have a valid password) [root]:
Enable Centralized Configuration Management (CCM) for this SBR node? [n]:
Configuring for use with generic database
Do you want to enable "Radius WatchDog" Process? [n]:
Do you want to enable LCI? [n]:
Do you want to configure for use with Oracle? [n]:
Removing oracle references from startup script
Do you want to configure for use with SIGTRAN? [n]:
Removing SIGTRAN references from startup script
Do you want to configure SNMP? [n]:
Configuring Admin GUI Webserver
Compatible Java version found :
Do you want to install custom SSL certificate for Admin WebServer? [n]:
Enable (e), disable (d), or preserve (p) autoboot scripts [e]:
```

```
The SBR Admin Web GUI can be launched using the following URL:
https://<servername>:2909
```

```
Configuration complete
```

```

-----
SBR 8.60.50006 cluster cambridge{0s,2sm,0m,4d}
on SunOS 5.10 Generic_141444-09 node wrx07(sm)is configured and processes are
down, may be reconfigured if desired
-----

```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

```
READY: last operation succeeded, node configured.
Enter the number of the desired configuration task or quit (4,q):
```

10. Enter **q** to quit.

11. Start the SSR process on the newly configured sm2 node:

- a. Execute:  
**./sbrd start ssr**
- b. Execute:  
**./sbrd status**
- c. Examine each line and ensure the SSR process is running without error.

12. Run **CreateDB.sh** script on sm2.

The purpose of running the **CreateDB.sh** script is to create certain files that are required to run the administrative shell scripts used to administer the session database in the cluster.

- a. Log in to sm2 as **hadm**.
- b. Navigate to the hadm user's home directory, **/opt/JNPRhadm/** by default.
- c. Execute:  
**CreateDB.sh**

13. Start the RADIUS process on sm2:

- a. Log in to sm2 as root.
- b. Execute:  
**./sbrd start radius**
- c. Execute:  
**./sbrd status**

The final cluster configuration looks as follows:

### Cluster Configuration

```
-----
[ndbd(NDB)]      4 node(s)
id=10   @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11   @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, starting, Nodegroup: 0)
id=12   @172.28.84.178  (mysql-5.7.25 ndb-7.6.9, starting, Nodegroup: 1)
id=13   @172.28.84.104  (mysql-5.7.25 ndb-7.6.9, starting, Nodegroup: 1)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1     @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2    @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)
```

```
[mysqld(API)]   4 node(s)
id=6     @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=7     @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)
id=58    @172.28.84.166  (mysql-5.7.25 ndb-7.6.9)
id=59    @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
```

- d. Examine each line and ensure the cluster is running with no errors.

## Non-Transition Server Method—Terminating Connections

If you can tolerate some downtime while the existing data nodes are stopped, the new configuration imposed, and all nodes restarted, that is the quickest and easiest method to incorporate the new data nodes. However, sessions are disconnected, and reconnection is not possible until all nodes come back online.



**CAUTION:** This procedure stops the entire cluster. You will not be able to process any requests from users.

To estimate how long this process takes, note the amount of time it takes to reconfigure one or two nodes.

Assuming the same basic configuration as in the previous examples of (0s), (2sm), 0(m), 2(d), the following procedure describes the high-level tasks involved in this method. Reference the previous procedures in this chapter for information about performing each task.

1. Stop the RADIUS processes on the sm1 and sm2 nodes.
2. Call **DestroyDB.sh** as user hadm on either the sm1 or sm2 node.
3. Stop the cluster on sm1.
4. Stop the SSR process on sm2.
5. Verify that the SSR processes are stopped on the two existing data nodes.
6. Install the SBR Carrier software on the two new data nodes in the expansion kit.
7. Run the configure script on sm1 using option 2 to update the cluster definition files.
8. Distribute the updated cluster definition files to all nodes including the two new data nodes.
9. Run the **clean** command on all four of the existing nodes (sm1, sm2, d1, and d2).
10. Start the SSR process on each node one at a time.
11. Run **CreateDB.sh** on the sm1 node.
12. After **CreateDB.sh** has finished running on sm1, repeat it on sm2 as user hadm.
13. Add the IP address pools and ranges using the administrative scripts.
14. Start the RADIUS processes on sm1 and sm2 one at a time.



# SIGTRAN Support for Steel-Belted Radius Carrier

---

Migrating Signalware | **310**

Installing Signalware 9 on Linux | **315**

Configuring SIGTRAN/IP Network Communication Files | **332**

---

# Migrating Signalware

## IN THIS CHAPTER

- [Migrating an Existing Signalware Configuration | 310](#)
- [Backing Up the Existing Signalware Configuration | 310](#)
- [Applying the Saved Signalware Configuration to the New Server | 312](#)

This chapter contains basic Signalware 9 backup and restore procedures, the tools that you can use to move an existing configuration to a new server to prepare for a new Steel-Belted Radius Carrier installation.

This chapter contains these topics:

## Migrating an Existing Signalware Configuration

If you have an existing Signalware installation, for instance one supporting Signalware 9 with Service Pack 6.3, you can migrate the Signalware configuration to a new server to support Steel-Belted Radius Carrier Release 8.6.0. Migrating your Signalware installation is a two-step operation: backing up your existing Signalware configuration and then restoring it to a new Signalware 9 installation on the server that you intend to support Steel-Belted Radius Carrier.

If you are installing on a server that does not already have a Signalware license, you need to update your license because each license is tied to a specific host.

## Backing Up the Existing Signalware Configuration

To back up the existing Signalware 9 configuration:



1. Retrieve your Signalware LSN number, which can be used to track the license.

To view your Signalware LSN number, execute:

```
$ cat ULCM_Signalware_9_license.lic | grep LSN
```

The system responds with the LSN number. Record it.

2. Back up your license file (`/etc/ULCM_Signalware_9_license.lic`) by copying it to a safe location.

The Signalware license is tied to the workstation by a unique host identification code, so if you use a different chassis as the new server, you must obtain a new or reissued license for that chassis.

You cannot physically transfer license files from one chassis to another, but you can save and restore them to a new installation on the same chassis.

3. Convert any **rc** (recently changed) files to **db** files using the MML **BACKUP-NODE** command.

- a. Make sure that Signalware is running.
- b. Change directories to `$OMNI_HOME/bin` and run the `swmml` utility.
- c. From within `swmml` run **BACKUP-NODE**.
- d. In another window, look in `$OMNI_HOME/CE NAME/dffile` for a file named **archive.xxx** that is time stamped with the date and time you ran the **BACKUP-NODE** command.
- e. Type **exit** to close `swmml`.

To learn more about **BACKUP-NODE**, at a system prompt, execute:

```
man BACKUP-NODE
```

4. Use the **DFcat** command to save the contents of the following files to a safe place such as a directory on the new server, a network file system, or a CD.

These files are in the `$OMNI_HOME/CENAME/dffile` directory:

- `cestart. $SHM`
- `archive.NODENAME. $SHM .date`
- `db.nodename.mtp. $SHM .pri`
- `db.nodename.sccp. $SHM .pri`

Where `$SHM` represents the value of the environment variable.

For example:

```
$OMNI_HOME/bin/DFcat cestart. $SHM >SAFE_PLACE/cestart. $SHM
```

## Applying the Saved Signalware Configuration to the New Server

After you install Signalware 9 (as described in [“Installing Signalware 9 on Linux” on page 315](#)), restore the saved configuration to that installation.

To restore the Signalware configuration:

1. Ensure that Signalware is not running. To stop Signalware, execute the **terminate** command:  
`$OMNI_HOME/bin/Terminate 0`
2. Copy or send through FTP the archived backup copies of **NODENAME.\$SHM.date** and **cestart.\$SHM** to **\$OMNI\_HOME/ulcm/bin** on the new server.
3. Create a new **process.mml** file from the old server's settings.
  - a. Execute **DFcat** on the **cestart.\$SHM** file from the old server to create **cestart.txt**.

For example, execute:

`$OMNI_HOME/bin/DFcat cestart.$SHM > cestart.txt`

Output is similar to this example:

```
.CE tssunblade100
GMT . 0 DEF DEF 0 0 /opt/JNPRsbr/radius/authGateway -name GMT -port 2000
-host tssunblade100 -node MGW -prot C7 -trace -conf
/opt/JNPRsbr/radius/conf/authGateway.con f -lri 1 -lpc 2002 -lssn 6 -appctx
3 -debug 0xff
@MGW C7M3UA
```

- b. Copy the **cestart.txt** file to create a new file named **process.mml**.

The new file should contain three lines, but **DFcat** might have inserted line breaks in the long string of the second line of the output that affect the appearance of the file. Delete these extra line breaks.

Depending on the string and your text editor, you may need to insert a space where the line breaks had been, to separate the items in the command string. After you delete the line breaks, make sure the command string entry is one long, unbroken, string.

4. Edit the new **process.mml** file.

This file contains the **START-PROCESS** command for the new server, built from the syntax used on the old server.

- a. Add a new first line to the file that contains the string:  
**START-PROCESS:**
  - b. Create a CE entry from the old server's CE data.

Change the line that reads (in the example):

**.CE tssunblade100**

to an attribute-and-field format that reads:

**CE="tssunblade100"**

- c. Create a NAME entry.

The first set of characters on the second line of the source file is the name of the server. Edit this line to change the string into a valid NAME entry (in the example) on its own line:

**GMT**

to the attribute-and-field format:

**NAME="GMT"**

- d. Create an EXEC entry.

Most of the remainder of the file contents is the invocation of the **authGateway** program, the command syntax the program uses when it starts. If the old server was functioning correctly, do not make any changes to the string syntax. You do need to clean up the raw string.

- Delete any characters at the beginning of the line that precede the command string (the command string begins with **/opt/JNPRsbr/radius...**).

In the example, the characters to delete are:

**. 0 DEF DEF 0 0**

- Insert an **EXEC="** field label before the command string.
- Insert a closing quotation mark at the end of the command string.

With those edits made, the file resembles this example:

```
START-PROCESS:
CE="tssunblade100"
NAME="GMT"
EXEC="/opt/JNPRsbr/radius/authGateway -name GMT -port 2000 -host tssunblade100
-node MGW -prot C7 -trace -conf /opt/JNPRsbr/radius/conf/authGateway.conf
-lri 1 -lpc 2002 -lssn 6 -appctx 3 -debug 0xff"
@MGW C7M3UA
```

- e. Delete the last line of the file, which contains the Signalware node name.

In the example, that line reads:

**@MGW C7M3UA**

- f. The file still contains both the old server's name and CE setting (**tssunblade100** and **GMT**). If the new server does not use those names, edit and update the **CE**, **NAME**, and **EXEC** fields with the new server settings.

- g. During the Signalware installation on the new server, if you changed the Signalware node name from the name used on the old server, edit and update that entry in the **EXEC** line (**-node MGW** in the example).
- h. Change the path to the **SBR Carrier authGateway** program from the 5.5.4 setting to the new default path for Steel-Belted Radius Carrier:  
**/opt/JNPRsbr/radius**
- i. Save and close the file.

5. Use the **DFcat** program on the archive file to create an **.mml** file.

For example:

```
$OMNI_HOME/bin/DFcat archive.NODENAME. $SHM .date > archive.mml
```

6. If the new Signalware installation does not inherit the same hostname and IP address used by the existing server, edit the **archive.mml** file to use the new server's information. You may change other settings at this time to match the new server environment.
7. Start Signalware.

Execute:

```
$OMNI_HOME/bin/go.omni
```

8. Load **archive.mml** first. Then load **process.mml**.

Execute:

```
$OMNI_HOME/bin/swmml -f archive.mml
```

```
$OMNI_HOME/bin/swmml -f process.mml
```

# Installing Signalware 9 on Linux

## IN THIS CHAPTER

- Operating System Prerequisite Packages | 316
- Other Prerequisites to Install Signalware on a Linux Platform | 318
- Installing Signalware | 319
- Troubleshooting Signalware Installation and Configuration | 328

Steel-Belted Radius Carrier supports Signalware installation on a Linux platform.

This chapter contains the basic procedure recommended by Juniper Networks for a new Signalware 9 with Service Pack 6 installation on a Linux platform. The installation example provided here is for your convenience and assumes that this is a clean first-time installation of Signalware.

Steel-Belted Radius Carrier supports a 64-bit kernel. Your Red Hat Enterprise Linux 7 operating system must comply with a 64-bit kernel with x86-based architecture. To verify the Linux kernel, run the following command:

```
[root@sbr-wf-eng1 ~]# cat /etc/redhat-release
Red Hat Enterprise Linux Server release 7.4 (Maipo)
[root@sbr-wf-eng1 ~]# uname -a
Linux sbr-wf-eng1 3.10.0-693.el7.x86_64 #1 SMP
Thu Jul 6 19:56:57 EDT 2017 x86_64 x86_64 x86_64 GNU/Linux
```

You must refer to the following Signalware documentation for instructions to install and configure Signalware on a Linux platform:

- *Signalware Read Me First*
- *Signalware Release Notes*
- *Signalware SP Release 9C Service Pack 6 Linux Installation Manual*

The Signalware documentation can be accessed from <https://my.ulticom.com>.

To troubleshoot errors that might occur during the installation and configuration process, see “[Troubleshooting Signalware Installation and Configuration](#)” on page 328.

The following topics are discussed in this chapter:

## Operating System Prerequisite Packages

This section lists the operating system packages that are required to install Signalware on a Linux platform. For more details about operating system requirements, see the *Signalware SP Release 9 Service Pack 6C Linux Installation Manual*.

**NOTE:** If you modify any kernel module, you need to execute the `swcommission(1s)` script again as well as configure Signalware.

To configure Signalware:

- Configure the computing element (`configurePlatform(1s)`)
- Commission the drivers (`swcommission(1s)`)
- Configure the nodes (`configureNodes(1s)`)
- Perform provisioning

### 64-Bit x86\_64 Red Hat Installations

On 64-bit Red Hat installations, both the 32-bit and 64-bit versions of Red Hat `glibc`, `glibc-devel`, `libstdc++`, and `libgcc` package RPMs must be installed.

To verify that the system contains both 32-bit and 64-bit versions of these GCC compiler tools packages, type:

```
rpm -q --queryformat "%{name}-%{version}-%{release}-%{arch}\n" glibc glibc-devel libstdc++ libgcc
```

This `rpm` command must return a report as follows where `xxx` is the version number:

- `glibc-xxx-x86_64`
- `glibc-xxx-i686`
- `glibc-devel-xxx-x86_64`
- `glibc-devel-xxx-i386`
- `libstdc++-xxx-x86_64`
- `libstdc++-xxx-i386`

- libgcc-xxx-x86\_64
- libgcc-xxx-i386

Only the 64-bit version of libstdc++-devel needs to be installed. To verify that the system contains the 64-bit version of the libstdc++-devel package, type:

**rpm -q --queryformat “{%name}-%{version}-%{release}-%{arch}\n” libstdc++-devel**

This must return a report as follows where xxx is the version number:

```
libstdc++-devel-xxx-x86_64
```

## Red Hat Enterprise Linux Release 7.3 and Above

Ensure that the following packages are installed on the Red Hat Enterprise Linux Release 7.3 operating system and above:

- kernel-3.10.0-229.11.1.el7.x86\_64 or higher, configured with smp enabled
- kernel-devel-3.10.0-229.11.1.el7.x86\_64 or higher, matching the running smp enabled kernel
- kernel-headers-3.10.0-229.11.1.el7.x86\_64 or higher, matching the running smp enabled kernel
- bash-4.2.46-12.el7 or higher
- bind-utils-9.9.4-18.el7 and bind-libs-lite-9.9.4-18.el7 or higher
- coreutils-8.22-12.el7 or higher
- rpm-4.11.1-25.el7 or higher
- gcc, version 4.8.3-9.el7 or higher
- 64 bit libstdc++, libstdc++-devel, version 4.8.3-9.el7 or higher
- 32 bit libstdc++, libstdc++-devel, version 4.8.3-9.el7 or higher
- 64 bit glibc, glibc-devel, version glibc-2.17-78.el7 or higher
- 32 bit glibc, glibc-devel, version glibc-2.17-78.el7 or higher
- xinetd-2.3.15-12.el7 or higher
- java-1.6.0-openjdk-1.6.0.0-1.20.b17 or higher
- perl-Env-1.04-2.el7

## Other Prerequisites to Install Signalware on a Linux Platform

The following are the prerequisites to install Signalware on a Linux platform:

- Uninstall any operating system Stream Control Transmission Protocol (SCTP) packages before installing Signalware. For example, none of the following packages should be installed:

```
[root@sbr-wf-eng1 ~]# yum info *sctp*
Loaded plugins: product-id, rhnplugin, subscription-manager
Updating Red Hat repositories.
Available Packages
Name           : lksctp-tools
Arch           : x86_64
Version        : 1.0.17
Release        : 2.el7
Size           : 227 k
Repo           : installed
From repo      : rhel-7-server-rpms
Summary        : User-space access to Linux Kernel SCTP
URL            : http://lksctp.sourceforge.net
License        : GPLv2 and GPLv2+ and LGPLv2 and MIT
Description    : This is the lksctp-tools package for Linux Kernel SCTP (Stream Control
Transmission Protocol) Reference Implementation. This package is intended to
supplement the Linux Kernel SCTP Reference Implementation now available in the
Linux kernel source tree in versions 2.5.36 and following. For more information
on LKSCTP see the package documentation README file, section titled "LKSCTP - Linux
Kernel SCTP." This package contains the base run-time library and command-line
tools.
```

- Copy the Signalware 9 license file to the /etc directory.
  - If you are migrating from an existing installation, you can use the existing Signalware 9 license as long as the license is applied on the same host (hostid).
  - To obtain a new license file for Signalware 9, contact your Juniper Networks representative.
- Create a UNIX user account called **siguser** as a member of the users group.
- Download Signalware packages to a temporary working directory or copy from a Signalware CD.
  - For Red Hat Enterprise Linux 7.3 or later, download "Signalware 9 SP6.v" (base package) and "Signalware 9 SP6.C" (upgrade).
- Install the Signalware packages.



- For Red Hat Enterprise Linux 7.3 or later, install the “Signalware 9 SP6.v” package first and then upgrade to “Signalware 9 SP6.C”.
- To troubleshoot errors that might occur during the installation and configuration process, see [“Troubleshooting Signalware Installation and Configuration” on page 328](#).

**NOTE:** On RHEL 7.5 or later, you must disable the kernel address space layout randomization (KASLR) feature on Linux Kernel. For more information on how to disable the KASLR feature, see [Table 35 on page 329](#).

## Installing Signalware

To install Signalware on a Linux platform:

1. Log in as root. We recommend that you not use a console port.
2. Stop the RADIUS process.
3. Copy the Signalware 9 license file to the **/etc** directory.
4. Create a UNIX group called **users** and verify that the group exists.
  - To add a group, execute the UNIX **groupadd** command:  
**\$ groupadd users**
  - To view a list of existing groups, execute the UNIX **cat** command:  
**\$ cat /etc/group**
5. Create a UNIX user account called **siguser** as a member of the **users** group.
6. Start the Signalware installation.

- For Red Hat Enterprise Linux 7.3 or later, get the “Signalware 9 SP6.v” package setup file (**swsetup**) by running `<package-directory>/rpm -Uvh --replacefiles OMNI-UTIL-K269.02SP6EECNv.OMNIUTIL.ECN9S6v.7-6.x86_64.rpm`.

`<package-directory>` is the directory path where the “Signalware 9 SP6.v” package setup file (**swsetup**) is located.

Now, you can find the **swsetup** file in the `/opt/swintall/bin` location.

Start the “Signalware 9 SP6.C” installation by executing `/opt/swintall/bin/swsetup`.

The script prompts you for a user identifier.

7. Enter the unique user that you had created earlier. For example, **siguser**.

The script prompts you for a scheduling priority.

8. Press Enter to accept the default value of 10.

The system checks for previous installations, a valid package file, and updates. Then the Product Menu is displayed.

```
HOST: bng-lnx-perf2          Ulticom (R) Product Menu          siguser (root)
                             Product Menu                     10 May 2018 12:15

1 = Signalware ..... Develop/Deploy SS7 Services
2 = nSignia ..... SS7 and IP Networks Convergence
3 = WLAN Libraries ..... Authentication/SMS Library Source
4 = Exit

>
Type 1-4 <enter>; <esc> or F11=Previous Menu; F12=Help; ?<enter>=Status
```

9. Enter 1 at the command prompt to install Signalware.

```
HOST: bng-lnx-perf2          Ulticom (R) Product Menu          siguser (root)
                             Signalware Main Menu             10 May 2018 12:16

1 = Install/Configure (Signalware is uninstalled or off-line)
2 = Online Upgrade (Signalware is installed and running)
3 = Installation Status and Reports
4 = Installation Maintenance
5 = Configuration Maintenance
6 = Start an Installed Instance of Signalware
```

```
>
Type 1-6 <enter>; <esc> or F11=Previous Menu; F12=Help; ?<enter>=Status
```

10. Enter 1 at the command prompt to install and configure Signalware.

The Install/Configure menu is displayed.

```
HOST: bng-lnx-perf2      Ulticom (R) Product Menu      siguser (root)
                        Install/Configure          10 May 2018 12:16

1 = [ ]Limit Installations to a Single Instance
2 = [X]Allow Multiple Installation Instances of Signalware
3 = Perform Initial Signalware Installation and Configuration

>
Type 1-3 <enter>; <esc> or F11=Previous Menu; F12=Help; ?<enter>=Status
```

11. Enter 3 at the command prompt to install and configure Signalware for the first time.

The Hosts menu appears, prompting you to enter hostnames in a cluster.

12. Press Enter because this is a single-node installation.

The Initial Install menu is displayed.

```
HOST: bng-lnx-perf2      Ulticom (R) Product Menu      siguser (root)
                        Initial Install          10 May 2018 12:17

1 = [ ]Install Packages
2 = [ ]Configure Platform
3 = [ ]Commission Instance
4 = [ ]Configure Nodes
5 = [ ]Start Signalware
6 = Done

>
Type 1-6 <enter>; <esc> or F11=Previous Menu; F12=Help; ?<enter>=Status
```

13. Enter 1 to select Install Packages.

The Select Instance screen is displayed and the script supplies the current (default) target directory.

For step-by-step Signalware installation procedures, see the *Signalware SP Release 9 Service Pack 6C Linux Installation Manual*.

14. Press Enter to accept the default target directory or specify an alternative location and press Enter.

The script prompts you to confirm the location.

15. Enter Y to continue.

The Package Directory screen is displayed.

16. Enter the full path to the directory containing the packages.

The packages are validated and installed.

The Protocol Selection menu is displayed.

17. Select the protocol that is appropriate to your location and system. Then press Enter.

As each protocol is activated, the screen is displayed again. When all protocols have been activated, type 0 and press Enter to continue.

**NOTE:** We recommend that you select only one protocol, typically option “2 = ITU Protocol”. Selecting multiple protocols may cause routing problems and complicate the MML configurations in later steps.

The protocol you select here must match the protocol choice you make in the later steps. For example, if you have selected option “2 = ITU Protocol”, then you must select the protocol “C7”.

The following information helps you to map the protocol selection:

- ANSI protocol corresponds to the A7 protocol
- ITU protocol corresponds to the C7 protocol
- CHINA protocol corresponds to the CH7 protocol
- JAPAN protocol corresponds to the J7 protocol

The Interface Selection menu is displayed.

18. Type 3 to select M3UA (SIGTRAN) (the interfaces supported by Steel-Belted Radius Carrier) and press Enter.

When the screen is redisplayed, type 0 and press Enter to continue.

The first page of the Select Packages menu is displayed.

19. Review the packages to be installed. (Generally, accept the defaults and do not deselect any packages.)

**NOTE:** The selected packages are based on the choices that you made earlier. We recommend that you do not deselect any packages.

20. Type 0 and press Enter to accept the default selections and continue.

A confirmation screen is displayed.

21. Enter Y to confirm the selection of all packages and install them.

Status messages are displayed while the packages are installed.

22. When Installation Complete is displayed, press Enter to continue.

The More Packages screen is displayed, and you are prompted to indicate whether you have completed installing packages.

23. Enter N to install the Signalware 9 SP6.C package.

24. Start the Signalware 9 SP6.C installation by entering the <package-directory> directory where the Signalware 9 SP6.C package is located.

The Installation Type menu is displayed.

```
1 = Quick Install
2 = Advanced Install
```

25. Enter 1 to select the Quick Install option.

- Quick Install option automatically updates any packages that have the prerequisites installed.
- Advanced Install option identifies the packages to be updated by querying the user.

When Installation Complete is displayed, press Enter to continue.

The Initial Install menu is displayed.

```
HOST: bng-lnx-perf2      Ulticom (R) Product Menu      siguser (root)
/opt/ulcm                Initial Install      10 May 2018 12:43

1 = [X]Install Packages
2 = [ ]Configure Platform
3 = [ ]Commission Instance
4 = [ ]Configure Nodes
```

```

5 = [ ]Start Signalware
6 = Done

>
Type 1-6 <enter>; <esc> or F11=Previous Menu; F12=Help; ?<enter>=Status

```

26. Enter 2 to configure the platform.

The Configure Platform menu is displayed.

```

HOST: bng-lnx-perf2          Ulticom (R) Product Menu          siguser (root)
/opt/ulcm                   Configure Platform              10 May 2018 12:44

1: Instance Install Directory.....: /opt/ulcm

>
Type Line Number to Change; Done=0 [0]

```

Enter **0** to accept the installation directory that you entered earlier.

After configuring the platform, the following prompt appears:

```
Press enter to continue:
```

27. Press Enter to continue with the remaining steps of the installation.

The Initial Install menu is displayed.

28. Enter 3 at the command prompt to commission Signalware.

This step installs all the Signalware drivers and starts the Signalware daemons.

The following is the sample output of a successful commission of Signalware:

```

Executing /opt/ulcm/conf/swcommission with OMNI_HOME=/opt/ulcm on CE sbr-wf-eng1
..
Executing swdecommission .....

```

```

Removing Signalware information from /etc/sysctl.conf
Removing Signalware information from /etc/ld.so.conf

```

```

Commissioning instance 1 in progress.....

```

```

.
installing sbin/sctpd
installing sbin/ulcm_start_native_sctp_wrap
installing bin/sctp_addr
installing bin/sctp_catcher
installing bin/sctp_discard
installing bin/sctp_echo
installing bin/sctp_pitcher
installing bin/sctp_stat
installing bin/sctp_monitor
installing bin/sctp_telnet
installing bin/sctp_daytime
installing lib/libsctp.so
installing lib/libsctp.a
installing include/sctp.h
installing include/sctp_notify.h.
Starting SCTP Daemon...
.....
SW_HOME=/opt/ulcm (instance 1) is commissioned.

```

After the Signalware is commissioned successfully, the following prompt is displayed:

```

Press enter to continue:

```

29. Press Enter to continue with the remaining steps of the installation.

The Initial Install menu is displayed.

30. Enter 4 at the command prompt to configure the node.

The Enter SHM screen is displayed.

```

HOST: bng-lnx-perf2      Ulticom (R) Product Menu      siguser (root)
/opt/ulcm                Enter SHM                        10 May 2018 12:51

```

```

In order to perform this operation, the system needs a SHM value.

```

```
>
Enter SHM to Configure []:
```

31. Enter the shared memory value based on the user ID of the siguser user that you created earlier (in Step 5).

The Signalware User screen is displayed.

32. Enter the username for running Signalware: siguser.

The Configure Nodes menu is displayed.

**NOTE:** When running configure nodes and selecting the protocol, it is recommended to select C7 protocol as it is the supported configuration. authGateway module does not support AC7 protocol. And it is not recommended to select AC7.

```
HOST: bng-lnx-perf2          Ulticom (R) Product Menu          siguser (root)
/opt/ulcm                   Configure Nodes              10 May 2018 12:54

1: Instance Install Directory.....: /opt/ulcm
2: Shared Memory.....: 225
3: Username.....: siguser

>
Type Line Number to Change; Done=0 [0]
```

33. Enter 0 to accept the values.

The Enter number of nodes screen is displayed.

34. We recommend that you accept the default value of 1.

The Enter node name screen is displayed.



35. Enter the name of the node.

The Select Protocol screen is displayed.

```

HOST: bng-lnx-perf2      Ulticom (R) Product Menu      siguser (root)
/opt/ulcm                Select Protocol              10 May 2018 12:54

1 = A7
2 = C7
3 = CH7
4 = J7

>2
Type 1-4 <enter>; <esc> or F11=Previous Menu; F12=Help; ?<enter>=Status

```

**NOTE:** The protocol you select here must match the corresponding protocol selected in the earlier steps. For example, if you have selected option “2 = ITU Protocol”, then you must select the protocol “C7”.

The following information helps you to map the protocol selection:

- A7 protocol corresponds to the ANSI protocol
- C7 protocol corresponds to the ITU protocol
- CH7 protocol corresponds to the CHINA protocol
- J7 protocol corresponds to the JAPAN protocol

The AC7 protocol (if it appears) corresponds to the hybrid ANSI/ITU. We recommend that you avoid selecting the AC7 protocol because it is not a proper superset of the A7 and C7 protocols. Selecting the AC7 protocol may cause routing problems and complicate the MML configurations in later steps.

36. Enter the protocol for the node.

The Select Protocol Type screen is displayed.

37. Enter the protocol type for the node. Select 3 for SIGTRAN.

The Select Options screen is displayed.

38. After the logical node and shared memory (SHM) values are configured successfully, the following prompt is displayed:

```
Press enter to continue:
```

39. Press Enter to continue with the remaining steps of the installation.

The Initial Install menu appears.

40. Enter 5 at the command prompt to start Signalware.

41. Enter 4 at the command prompt to exit the Initial Install menu.

This completes the Signalware installation on a Linux platform.

#### NOTE:

The following points have to be considered, if you are planning for Installing/Upgrade of Signalware version to 9s6C on RHEL 7.6 or Later version.

- Before installing Signalware on RHEL 7.6 or later, you must disable the Hardened User-copy feature, which is enabled by default on RHEL 7.6 or later.
- Signalware is incompatible with Hardened User-copy feature and the feature MUST be disabled to ensure the Omnimon debugger is 100% safe to run in production network.
- For detailed information of system prerequisites and installing procedure, refer [Signalware Installation/Upgrade](#).

## Troubleshooting Signalware Installation and Configuration

This section describes how to troubleshoot errors that might occur during the installation and configuration of Signalware on a Linux platform.

Table 35: Troubleshooting Signalware Installation and Configuration

Error	Solution
<p>  ERROR  </p> <p>Cannot find the kernel source config file /proc/config.gz or /lib/modules/3.10.0-693.el7.x86_64/build/.config. Please ensure that a config file that matches the running kernel exists.</p> <p>rm: cannot remove `.`: Is a directory rm: cannot remove `..`: Is a directory</p>	

Table 35: Troubleshooting Signalware Installation and Configuration (*continued*)

Error	Solution
	<p>This error occurs when you configure the platform if the running kernel != kernel-devel version.</p> <p>If you have used a precompiled kernel, you must make sure that your kernel sources have the same kernel configuration used for compiling the kernel. You must copy the kernel configuration from /proc/config.gz, or /boot/config-&lt;version&gt; into /lib/modules/uname -r/build/.config by using the gunzip &lt; /proc/config.gz &gt; /lib/modules/\$(uname -r)/build/.config command.</p> <p>Recompile and install ndiswrapper.</p> <pre>[root@sbr-wf-eng1]# yum info kernel more Loaded plugins: product-id, rhnplugin, subscription-manager Updating Red Hat repositories. Installed Packages Name           : kernel Arch           : x86_64 Version        : 3.10.0 Release        : 693.el7  [root@sbr-wf-eng1]# yum info kernel-devel more Loaded plugins: product-id, rhnplugin, subscription-manager Updating Red Hat repositories. Installed Packages Name           : kernel-devel Arch           : x86_64 Version        : 3.10.0 Release        : 693.el7  [root@sbr-wf-eng1]# yum info kernel-headers more Loaded plugins: product-id, rhnplugin, subscription-manager Updating Red Hat repositories. Installed Packages Name           : kernel-headers Arch           : x86_64 Version        : 3.10.0 Release        : 693.el7  [root@sbr-wf-eng1]# yum install kernel-devel-3.10.0-693.el7 [root@sbr-wf-eng1]# yum install</pre>

Table 35: Troubleshooting Signalware Installation and Configuration (*continued*)

Error	Solution
	kernel-headers-3.10.0-693.el7
<p>The following command did not complete successfully.</p> <p>Command: /opt/ulcm/conf/swcommission on CE sbr-wf-eng1</p> <p>Error Code: 1</p> <p>Error Text:</p> <p>Signalware swdecommission Script Error. inetd is NOT RUNNING!!</p> <p>&gt; Do you wish to see the complete error output (Y/N)[Y]:</p>	<p>This error occurs when you commission the Signalware 9SP6.0 without upgrading to Signalware 9SP6.C</p> <p>Perform the following steps:</p> <ol style="list-style-type: none"> <li>1. At the <b>Do you wish to see the complete error output (Y/N):</b> prompt, enter N.</li> <li>2. Enter 6 -override to skip the remaining steps.</li> <li>3. Press F11 to return to the previous menu. The Initial Install menu appears.</li> <li>4. Enter 1 at the command prompt to install the Signalware 9SP6.A package.</li> </ol>
<p>RHEL 7.5 defaults to enabling kernel address space layout randomization (KASLR). This causes Signalware to crash when loading the ulcm_kmem module into the kernel because this module is not KASLR aware and expects the kernel memory address space to be consistent.</p> <p>To find out more about KASLR, refer to the RHEL 7.5 kernel release notes.</p>	<p>On RHEL 7.5 or later, you must set the nokaslr kernel command line option via the bootloader to force a boot to a non-KASLR kernel.</p> <p>Perform the following steps:</p> <ol style="list-style-type: none"> <li>1. Edit the <b>GRUB_CMDLINE_LINUX</b> key in the <b>/etc/default/grub</b> file to add the new parameter <b>nokaslr</b> as in the following example: <pre>GRUB_CMDLINE_LINUX="net.ifnames=0 biosdevname=0 crashkernel=auto rd.lvm.lv=vg_vda/root console=ttyS0,115200n8 console=tty0 nokaslr"</pre> </li> <li>2. Run the <b>grub2-mkconfig -o /boot/grub2/grub.cfg</b> command.</li> <li>3. Reboot the server. This disables the KASLR feature on Linux Kernel.</li> <li>4. Perform Signalware swcommission.</li> <li>5. Start Signalware.</li> </ol>

# Configuring SIGTRAN/IP Network Communication Files

## IN THIS CHAPTER

- Starting and Stopping Signalware | 332
- Configuring Communication Files Overview | 333
- Communication Pathways and Corresponding Files | 334
- Configuration Activities | 335
- Signalware MML Commands | 335
- Defining Links, Link Sets, and Route Sets | 336
- Configuring authGateway and GWrelay Applications for HLR Communication | 338
- Loading the MML Configuration Settings | 354
- Sample authGateway Command and File | 355

This chapter discusses configuring Signalware to communicate through SIGTRAN (provided as a module of Signalware) to allow SS7 communication over IP networks. These topics are included in this chapter:

## Starting and Stopping Signalware

To start the Signalware system manually, execute: `/$OMNI_HOME/bin/go.omni`.

To stop Signalware, execute: `/$OMNI_HOME/bin/Terminate 0`.

## Configuring Signalware to Start Automatically on Reboot

We recommend that you configure Signalware so that it restarts automatically when the system is rebooted. If you choose not to configure Signalware in this way, you need to restart Signalware with the `go.omni` command whenever the system reboots.

To configure Signalware to start automatically on reboot:

1. Change to superuser mode.

Execute:

```
$su
```

2. Copy the **S91omni** file as **omni** into the directory **/etc/init.d**. The **S91omni** file is located at **radius/samples/SIM**.

3. Edit the **omni** file to change the **SHM=** line to specify the SHM value set during Signalware installation.

Example:

```
SHM = 102
```

To find your user ID (which is the same as the SHM), execute **id** at the prompt.

4. Change the run permissions.

Execute:

```
$ chmod 755 omni
```

5. Create links into the **/etc/rc3.d** directory.

Execute:

```
$ ln -s /etc/init.d/omni /etc/rc3.d/S91omni
```

6. Make sure a subdirectory called **Logs** is in the directory where Signalware is installed. If necessary, create the **Logs** directory.

## Configuring Communication Files Overview

When configuring the communication pathways between Steel-Belted Radius Carrier and equipment in the SS7 network, you need to configure certain files based on your choices of the following:

- Type of network equipment Steel-Belted Radius Carrier communicates with for processing authorization requests. (See [Table 36 on page 333](#) for information of which type of equipment is used.)
  - HLR
- Type of network.
  - SS7 over IP (SIGTRAN)

Table 36: Network Equipment Used for Authorization

Action Needed to Process Access-Request	Network Equipment
Obtain SIM triplets	HLR

Table 36: Network Equipment Used for Authorization (*continued*)

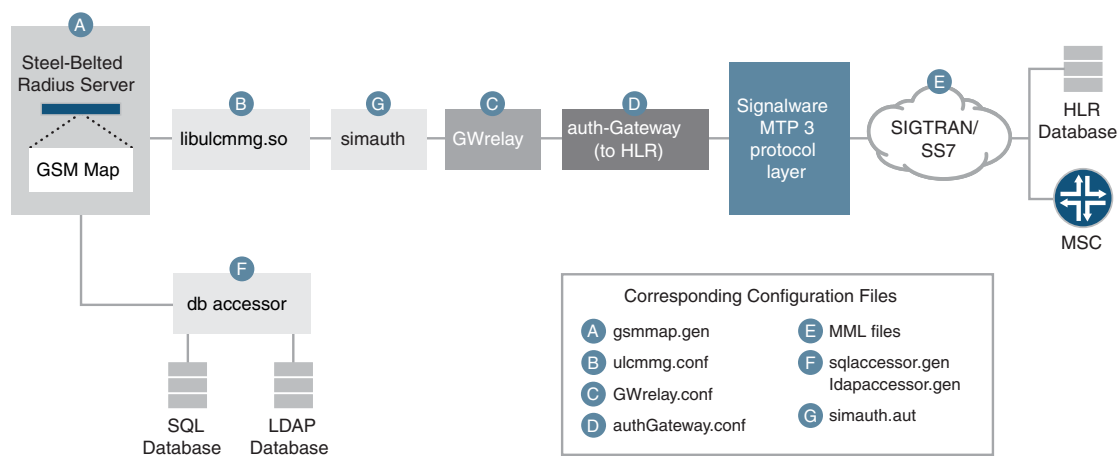
Action Needed to Process Access-Request	Network Equipment
Obtain AKA quintets	HLR
Obtain IMSI (given the MSISDN)	HLR
Obtain MSISDN (given the IMSI)	HLR

**NOTE:** If quintets are received but triplets are needed, then Steel-Belted Radius Carrier converts the quintets to triplets according to specification 3G TS 33.102, available at <http://www.3gpp.org>. Remember to configure and support MAPv3.

## Communication Pathways and Corresponding Files

Figure 15 on page 334 shows the communication pathways between the Steel-Belted Radius Carrier server and the network equipment. It also lists the files that need to be configured for each segment in the communication pathways.

Figure 15: Steel-Belted Radius Carrier Libraries and Corresponding Configuration Files





## Configuration Activities

To configure the communication pathways:

1. Install Signalware. For details on the recommended procedure to install and run Signalware 9, see [“Installing Signalware 9 on Linux” on page 315](#).
2. Define links, link sets, and route sets with MML commands for SS7 or SIGTRAN. See [“Defining Links, Link Sets, and Route Sets” on page 336](#).
3. Configure the authGateway and GWrelay applications for HLR communication. See [“Configuring authGateway and GWrelay Applications for HLR Communication” on page 338](#).
4. Load the MML configuration settings. See [“Loading the MML Configuration Settings” on page 354](#).
5. Start Signalware.

## Signalware MML Commands

After Signalware is installed, you can configure and provision Signalware using commands sent to the Signalware system. For details about installing and running Signalware, see [“Installing Signalware 9 on Linux” on page 315](#).

These commands are in Man-Machine Language (MML). You can input MML commands individually using the SWMML program, or save them in a file. The procedures in this chapter assume that you save the MML commands to `.mml` text files and execute them as described in [“Loading the MML Configuration Settings” on page 354](#).

The basic activities that require MML commands are:

- Setting up link sets, links, and routes—See [“Defining Links, Link Sets, and Route Sets” on page 336](#).
- Configuring the authGateway location and startup information—See [“Configuring authGateway and GWrelay Applications for HLR Communication” on page 338](#).
- Loading the MML configuration settings—See [“Loading the MML Configuration Settings” on page 354](#).

To view a list of all MML commands and definitions, enter: **man MML\_Intro**

To view specific information about any MML command, enter: **man cmdname**.

For example:

```
$ man CRTE-LSET
```

## Defining Links, Link Sets, and Route Sets

Links identify point-to-point connections of an adjacent node. Link sets are sets of parallel links that can be used for load sharing. Routes identify the final node destination. (The order in which they are defined is important.) [Figure 16 on page 337](#) provides a simple example of links, link sets, and route sets. In the case of a SIGTRAN installation, the link set most likely consists of a single link because redundancy might not be needed.

You use MML commands to set up links, link sets, and routes to identify the path between Steel-Belted Radius Carrier, the adjacent link, and the final destination (HLR or MSC).

Setting up links, link sets, and route sets involves the MML commands listed in [Table 37 on page 336](#). For more information about the syntax and usage of the MML commands, see [“Signalware MML Commands” on page 335](#).

**Table 37: MML Commands for Defining Links, Link Sets, and Route Sets**

Action	MML Command for SS7 Networks	MML Command for SIGTRAN Networks
Identify the local point code and network indicator (NI)	CRTE-OSPC	CRTE-OSPC
Create a link set and assign it a point code	CRTE-LSET	CREATE-M3UA-LSET
Create one or more links that belong to the link set	CRTE-SLK	CREATE-M3UA-SLK
Create a route set that identifies the final destination	CRTE-RSET	CRTE-M3UA-RKEY
Allow the route set to be used	ALW-RSET	ALW-RSET
Activate the links	ACTV-SLK	ACTV-SLK

### Example MML Commands

The examples in this section illustrate the MML commands used to create links, link sets, and route sets.

Figure 16: Links, Link Sets, and Route Sets

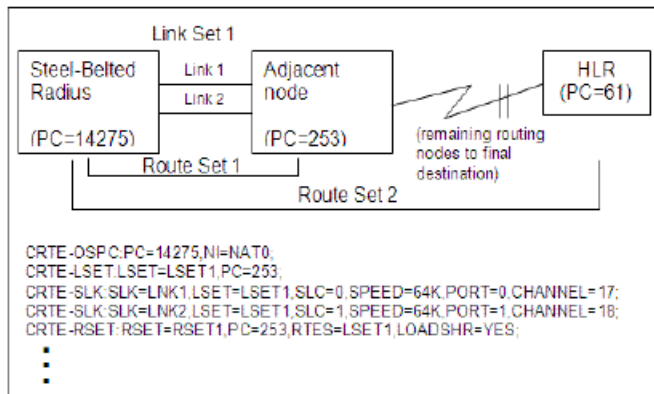
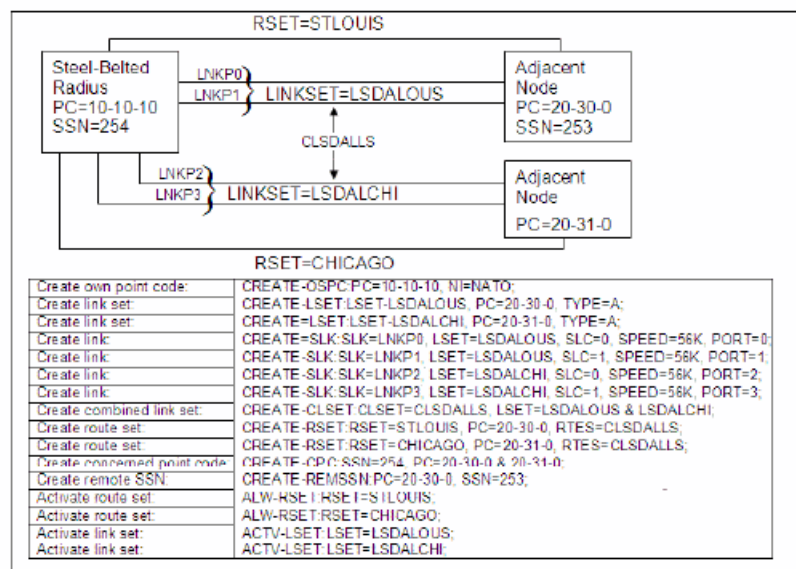


Figure 17: MML Provisioning

**SS7 Example: Creating Links, Link Sets, and Route Sets**

```

CRTE-OSPC:PC=14275,NI=NAT0;
CRTE-LSET:LSET=LSET1,PC=253;
CRTE-SLK:SLK=LNK1,LSET=LSET1,SLC=0,SPEED=64K,PORT=0,CHANNEL=17
CRTE-RSET:RSET=RSET1,PC=253,RTES=LSET1,LOADSHR=YES;
CRTE-RSET:RSET=RSET2,PC=61,RTES=LSET1,LOADSHR=YES;
ALW-RSET:RSET=RSET1;
ALW-RSET:RSET=RSET2;
ACTV-SLK:SLK=LNK1;

```

**SIGTRAN Example: Creating Links, Link Sets, and Route Sets**

```

CREATE-OSPC:PC=4004,NI=INT0;

```

```

CREATE-M3UA-LSET:LSET=IPLSET,TYPE=IPSP-IPSP, RADDR=207.46.20.60,PC=5005;
CREATE-M3UA-SLK:SLK=IPSLK,LSET=IPLSET,LADDR=72.5.124.61,RADDR=207.46.20
.60,MODE=CONNECT,LPORT=2906;
ACTIVATE-M3UA-SLK:SLK=IPSLK;

CREATE-RSET:RSET=IPRSET,PC=5005,RTES=IPLSET;
ALLOW-RSET:RSET=IPRSET;

CREATE-M3UA-RKEY:RKEY=RK1,TYPE=STATIC-AS,TRAFFIC-MODE=LOADSHARE,LSET=IP
LSET,DPC=4004,SI=SCCP,SSN=1&251&252;
ACTIVATE-M3UA-RKEY:RKEY=RK1;

```

## Configuring authGateway and GWrelay Applications for HLR Communication

The authGateway application manages all communication between SBR Carrier and the HLR. The authGateway application also implements the MAP (Mobile Application Port) protocol and MAP messages that get sent through the Signalware protocol stack and out to the HLR and back. Multiple authGateway instances can be used to process multiple authentication and authorization requests simultaneously. The GWrelay application is used to pass authentication requests between SBR Carrier and the authGateway instances in a round-robin method. The GWrelay application establishes an SCTP connection with each authGateway instance through unique source and destination ports.

Configuration of authGateway and GWrelay applications requires the activities described in the following sections of this chapter:

- [“Configuring the authGateway Routing Location Information” on page 338](#)
- [“Configuring the authGateway.conf File” on page 339](#)
- [“Configuring the authGateway Startup Information” on page 348](#)
- [Configuring the GWrelay.conf File on page 353](#)
- [Starting the GWrelay Process on page 353](#)
- [“Configuring the ulcmmg.conf File” on page 354](#)

### Configuring the authGateway Routing Location Information

This activity assigns the local routing options and the remote routing options using the MML commands listed in [Table 38 on page 339](#).

Table 38: MML Commands for Configuring authGateway Routing Location

MML Command	Description
CREATE-CPC	Identify the concerned point code (CPC), which is the destination point code and the local application (authGateway).
CREATE-REMSSN	Identify the point code of the HLR and the remote application.
CREATE-GT	Create a global title translation for the remote HLR (if Global Title routing is used).

For more information about the syntax and usage of the MML commands, see [“Signalware MML Commands” on page 335](#).

#### **Example 1—Global Title Routing Using Global Title Identification**

In the following example, these actions take place:

Line 1 (**CREATE-GT**): Global Title type translation is used so that digits 22201 are sent to PC, SSN=61, 6. RI=GT notifies 61,6 that it needs to find the next routing hop for the request.

**CREATE-GT:TT=0,NP=ISDN-MOB,NA=INT,DIG="22201",PC=61,SSN=6,RI=GT;**

#### **Example 2—PCSSN Routing Using Point Code Identification**

In the following example, these actions take place:

Line 1 (**CREATE-CPC**): authGateway is assigned a subsystem number (SSN) of 7 on the local host and the concerned point code on the HLR is identified as 61.

Line 2 (**CREATE-REMSSN**): The subsystem number (application) on the remote host is identified as 6.

Line 3 (**CREATE-GT**): Global Title type translation is used so that digits 22201 are sent to PC, SSN=61, 6. RI=PCSSN indicates that digits 22201 are handled by PC, SSN=61, 6.

**CREATE-CPC:PC=61,SSN=7;**

**CREATE-REMSSN:PC=61,SSN=6;**

**CREATE-GT:TT=0,NP=ISDN-MOB,NA=INT,DIG="22201",PC=61,SSN=6,RI=PCSSN;**

**NOTE:** MML commands are saved in MML files, which can be loaded into Signalware. See [“Loading the MML Configuration Settings” on page 354](#).

## **Configuring the authGateway.conf File**

The **authGateway.conf** file configures the following authGateway options:

- Remote routing options control how the remote HLR is addressed based on the incoming IMSI.
- Authorization options control whether or not a subscriber requesting an account is authorized for WLAN access, and which Steel-Belted Radius Carrier profile or native user is used.
- The **FetchMSISDNRoutingInfoLCS** parameter specifies the type of message that is used to request MSISDN information from an HLR or HSS, and the **CamelSupportedPhases** parameter specifies which CAMEL phase services are supported in the network.
- Common configurations of the authGateway process.
- Process-specific options specify settings related to the authGateway process.

### **[Routing-Configuration] Section**

Each line in the **authGateway.conf** file represents a target HLR, where each HLR has its own routing options and authorization options. Indicate each HLR listed in this file with the initial digits of the subscriber password, specified by the **odigits** option.

Table 39 on page 340 lists the remote routing options for the **authGateway.conf** file.

**Table 39: authGateway.conf [Routing-Configuration] Section Syntax**

Option	Purpose
bs	Bearer Service. See <a href="#">“Authorization Options” on page 341</a> .
msisdn	The <b>msisdn</b> option can be used in place of <b>ndigits</b> and <b>odigits</b> when no translation is required. See <a href="#">“Example 2—authGateway.conf file [Routing-Configuration] Section” on page 343</a> .
ndigits	Replacement digits for numbering plan translation (hybrid IMSI).
odb	Operator-Determined Barring. See <a href="#">“Authorization Options” on page 341</a> .
odigits	<p>Initial digits of IMSI or password for this HLR. For each request, the first digits of the IMSI are compared with <b>odigits</b>. The first line of the configuration file that matches is selected for the current request.</p> <p>If the routing indicator (<b>rr</b>) is 0 (Global Title), the leading digits are replaced with the new digits (<b>ndigits</b>) to perform the numbering plan translation.</p> <p>Example of direct replacement:</p> <p>If the rule is “<b>odigits 12345 ndigits 98765</b>” and the IMSI is 123456789012345, the resulting digits are 987656789012345.</p> <p>Example of wildcard replacement:</p> <p>If the rule is “<b>odigits 12345* ndigits 98765</b>” and the IMSI is 123456789012345, the resulting digits are 98765.</p>

Table 39: authGateway.conf [Routing-Configuration] Section Syntax (*continued*)

Option	Purpose
rgti	(Global Title only) GTI value. 4 for C7; 2 for A7. (Usually 4.)
rnai	(Global Title only) Nature of Address Indicator.
rnpi	(Global Title only) Numbering Plan.  Acceptable values are:  1—ISDN/Telephony  3—DATA  4—TELEX  5—Maritime Mobile  6—Land/Mobile  7—ISDN/Mobile  10—British Telecom special 1  11—British Telecom special 2  14—Private Network
rpc	Remote Point Code. Point Code of HLR or MSC.
rri	Routing indicator - 0 for GT (Global Title), 1 for PC/SSN (Point Code/Subsystem Number).
rssn	Subsystem Number of HLR.
rtt	(Global Title only) Translation Type (usually 0).
ts	Teleservice. See <a href="#">“Authorization Options” on page 341</a> .

### Authorization Options

The HLR database includes authorization information that is assigned to each subscriber. Three authorization designations are relevant to Steel-Belted Radius Carrier with the SIM authentication module:

- BS (Bearer Service)
- TS (Teleservice)
- ODB (Operator-Determined Barring)

You can specify subscriber HLR authorization (and barred service) designations in the MAP Gateway **authGateway.conf** file.

**NOTE:** You can disable authorization completely from EAP-SIM (not fetch subscriber profile information from the HLR and not perform a SQL/LDAP query). For instructions about disabling authorization, see *“Disabling Authorization from EAP-SIM”* in the section on *Configuring the gsmmap.gen File for the SIM Authentication Module*, in the *SBR Carrier Reference Guide*.

Each line in the **authGateway.conf** file corresponds to an HLR in your network. Each line also specifies all potential authorization (and barred service) settings for any subscribers on this HLR.

Steel-Belted Radius Carrier with the SIM authentication module uses the service authorization information that you list for each HLR in **authGateway.conf**:

- When a TS or BS designation is assigned to a subscriber entry in the HLR database, Steel-Belted Radius Carrier with the SIM authentication module allows the subscriber the designated class of WLAN service upon authorization request.
- When an ODB designation is assigned to a subscriber, Steel-Belted Radius Carrier with the SIM authentication module denies the subscriber WLAN service upon authorization request.
- When you do not specify service designations for a HLR listed in **authGateway.conf**, then all subscribers on that HLR are authorized for WLAN service.
- You can specify up to six authorization strings of each type (TS, BS, or ODB) on any given line of **authGateway.conf**.

You can specify the service designations in **authGateway.conf**:

**bs n1: auth1**

**ts n2:auth2**

**odb n3:auth3**

Here, **ni** (**i=1,2,3**) is a decimal integer that specifies the setting, and **authi** (**i=1,2,3**) is the string returned from the MAP Gateway to Steel-Belted Radius Carrier with the SIM authentication module.

For example, you might specify the potential subscriber designations on one HLR with the following text in **authGateway.conf**:

**bs 26:B1A ts 33:TS21 odb 128:bar**



**NOTE:** If you require any HLR authorization strings to define different classes of service for your subscribers, you must also specify those TS, BS, and ODB authorization strings in certain files associated with the SIM authentication module. For information about how to match these strings to Steel-Belted Radius Carrier variables, see the “*simauth.aut [ProfileMap] Section*” of *Configuring EAP-SIM and EAP-AKA for the SIM Authentication Module* in the *SBR Carrier Reference Guide*.

**Example 1—authGateway.conf file [Routing-Configuration] Section**

(Lines are wrapped.)

```
odigits 2310 ndigits 2324 rnai 4 rnp 7 rgti 4 rtt 0 rri 0 rpc 3003 rsn 251 bs 12:gold bs 23:silver ts 91:bronze
ts 92:red ts 93:green odb 1:black aqua
```

```
odigits 31026 ndigits 32476 rnai 4 rnp 7 rgti 4 rtt 0 rri 1 rpc 3003 rsn 253 bs 23:morning bs 24:afternoon
ts 1:night
```

**Example 2—authGateway.conf file [Routing-Configuration] Section**

In this global title example, odigits and ndigits are the same and do not require translation. You can use the **msisdn** option in place of ndigits and odigits when no translation is required.

(Lines are wrapped.)

```
msisdn 31026 rnai 4 rnp 7 rgti 4 rtt 0 rri 0 rpc 3003 rsn 251 bs 12:gold bs 23:silver ts 91:bronze ts 92:red
ts 93:green odb 1:black aqua
```

**[Supported-MAP-Messages] Section**

The [Supported-MAP-Messages] section ([Table 40 on page 344](#)) of the **authGateway.conf** file specifies the method of fetching MSISDN and which CAMEL phase services are supported in the network.

Table 40: authGateway.conf [Supported-MAP-Messages] Syntax

Parameter	Description
FetchMSISDNRoutingInfoLCS	<p>Specifies the type of message that is used to fetch MSISDN information from an HLR or HSS.</p> <p>MSISDN information is usually fetched from an HLR through the d interface using the RestoreData message. Setting the <b>FetchMSISDNRoutingInfoLCS</b> parameter to 0 configures the authGateway process to interact with the HLR through the RestoreData message. The default SSN configured when the authGateway process starts is used as the originating SSN in the RestoreData message.</p> <p>MSISDN information is usually fetched from an HLR or HSS through the SLh or Lh interface using the SendRoutingInfoForLCS message. Setting the <b>FetchMSISDNRoutingInfoLCS</b> parameter to 1 configures the authGateway process to interact with the HLR or HSS through the SendRoutingInfoForLCS message. Because SBR Carrier acts as a GMLC in this case, the GMLC SSN (i.e. 145) is used as the originating SSN.</p> <p>By default, this parameter is set to 0.</p>
CamelSupportedPhases	<p>Specifies which CAMEL phase services are supported in the network.</p> <ul style="list-style-type: none"> <li>• If set to 0 or commented out, the RestoreData message populates only the mandatory parameter IMSI.</li> <li>• If set to 1, the network supports CAMEL phase 1 services.</li> <li>• If set to 2, the network supports CAMEL phase 2 services.</li> <li>• If set to 3, the network supports CAMEL phase 3 services.</li> <li>• If set to 4, the network supports CAMEL phase 4 services.</li> </ul> <p>By default, this parameter is set to 0.</p>

**[Common-AGW-Configurations] Section**

The [Common-AGW-Configurations] section ([Table 41 on page 344](#)) of the **authGateway.conf** file specifies common configurations for the authGateway process.

Table 41: authGateway.conf [Common-AGW-Configurations] Syntax

Option	Description
appctx	<p>Specifies the MAP protocol revision (2 or 3). Only MAPv3 retrieves quintets, so it must be used to support EAP-AKA.</p> <p>Default value is 2.</p>

Table 41: authGateway.conf [Common-AGW-Configurations] Syntax (continued)

Option	Description
connretry	Specifies the number of connection attempts.  Default value is 10.
conntimeout	Specifies the connection timeout in minutes.  Default value is 0 minute.
host	Specifies the local hostname. You must use the hostname associated with the IP address that the authGateway listen on. Also, you must ensure that the entry is coordinated with the <b>radius/GWrelay.conf</b> file (if authGateway is running as multiple instances) and <b>radius/conf/ulcmng.conf</b> (if authGateway is running as a single instance) files. If a hostname is not specified, 0.0.0.0 is used.
invkretry	Specifies the number of invoke retries.  Default value is 1.
invktimeout	Specifies the duration of invoke timeout in seconds.  Default value is 30 seconds.
lgti	(Global Title only) Specifies the local GTI value, usually 4 for C7 and 2 for A7.  Default value is 4.
lmsisdn	(Global Title only) Specifies the MSISDN value of the local node.  Default value is 0.
lnai	(Global Title only) Specifies the scope of the address value, such as whether it is an international number (includes country code) or a national number (no country code). <ul style="list-style-type: none"> <li>• 1—Subscriber number with no area code (example: 5551234)</li> <li>• 2—Unused</li> <li>• 3—National significant number with no country code (example: 2015551234)</li> <li>• 4—International number with country code (example: 12015551234)</li> </ul> Default value is 4.

Table 41: authGateway.conf [Common-AGW-Configurations] Syntax (*continued*)

Option	Description
lnp	<p>(Global Title only) Specifies the local numbering plan. Acceptable values are:</p> <ul style="list-style-type: none"> <li>• 1—ISDN/Telephony</li> <li>• 3—DATA</li> <li>• 4—TELEX</li> <li>• 5—Maritime Mobile</li> <li>• 6—Land/Mobile</li> <li>• 7—ISDN/Mobile</li> <li>• 10—British Telecom special 1</li> <li>• 11—British Telecom special 2</li> <li>• 14—Private Network</li> </ul> <p>Default value is 1.</p>
lpc	<p>Specifies the Local Point Code (PC).</p> <p>Default value is 0.</p>
lri	<p>Specifies the routing indicator used to address messages to the local node.</p> <ul style="list-style-type: none"> <li>• 0—Global Title</li> <li>• 1—PC/SSN</li> </ul> <p>Default value is 1.</p>
ltt	<p>(Global Title only) Specifies local translation type.</p> <p>Default value is 0.</p>
max_requests	<p>Specifies the maximum number of concurrent requests that can be handled by the authGateway process.</p> <p>Default value is 1000.</p>
monitor	<p>Activates the message activity monitor.</p>
no_rst	<p>Disables automatic restart of the authGateway process.</p>
node	<p>Specifies the node name.</p>

### [Process<name>] Section

The [Process<name>] section ([Table 42 on page 347](#)) of the **authGateway.conf** file contains the parameters that control authGateway process specific behavior.

**Table 42: authGateway.conf [Process<name>] Syntax**

Option	Description
debug	Specifies a debug level. Default value is 0.  <b>NOTE:</b> This parameter is reloaded whenever SBR Carrier receives a SIGHUP (1) signal.
lssn	Specifies the local subsystem number.  Default value is 7.
port	Specifies the port number used by the SCTP association with the client.
prot	Specifies the variant used (C7, A7, or CH7).
trace	Enables debug tracing and displays the trace information about the console. (Consists of a trace of all MAP messages that are formatted and sent down the stack.)  <b>NOTE:</b> We recommend setting this parameter to 0.  <b>NOTE:</b> This parameter is reloaded whenever SBR Carrier receives a SIGHUP (1) signal.
tracefile	Captures the trace information to a file. The filename follows the <b>-tracefile</b> switch. Include the directory in the filename.  <b>NOTE:</b> This parameter is reloaded whenever SBR Carrier receives a SIGHUP (1) signal.

### Example

If you are using two authGateway processes—for example, **GMT1** and **GMT2**, then two sections **[ProcessGMT1]** and **[ProcessGMT2]** must be added to the **authGateway.conf** file for the authGateway processes to startup. The following example explains this configuration:

```
[ProcessGMT1]

#Remote port specified in ulcmmg.conf
#Port number used by the SCTP association with the client
port=2003

#Variant used (C7, A7 or CH7)
prot=C7
```

```

#Enables Signalware library debug logging. Sets a debug level.
debug=1

#This enables debug tracing and displays the trace information about the console.
#Consists of a trace of all MAP messages that are formatted and sent down the stack.
#Use the tracefile option to capture the trace information to a file
trace=1

#Captures the trace information to a file. The filename follows the tracefile
#switch. Include the directory in the filename
tracefile=/opt/JNPRsbr/radius/conf/Trace1.out

[ProcessGMT2]

#Remote port specified in ulcmmg.conf
#Port number used by the SCTP association with the client
port=2005

#Variant used (C7, A7 or CH7)
prot=C7

#Enables Signalware library debug logging. Sets a debug level.
debug=1

#This enables debug tracing and displays the trace information about the console.
#Consists of a trace of all MAP messages that are formatted and sent down the stack.
#Use the tracefile option to capture the trace information to a file
trace=1

#Captures the trace information to a file. The filename follows the tracefile
#switch. Include the directory in the filename
tracefile=/opt/JNPRsbr/radius/conf/Trace2.out

```

## Configuring the authGateway Startup Information

The **CREATE-PROCESS** and **START-PROCESS** MML commands start the authGateway (by calling **authGateway.conf**), using options that you specify.

[Table 43 on page 349](#) describes the MML commands needed to configure and start authGateway.

Table 43: MML Commands for Configuring the Start of authGateway

MML Command	Description
<b>CREATE-PROCESS</b>	Identify the authGateway configuration file and the authGateway options.
<b>START-PROCESS</b>	Start the process.

For more information about the syntax and usage of the MML commands, see [“Signalware MML Commands” on page 335](#). See [“Loading the MML Configuration Settings” on page 354](#) for information about executing the MML commands.

[Table 44 on page 349](#) lists the mandatory MML options to be used with the **CREATE-PROCESS** command.

Table 44: authGateway Process Options Used with CREATE-PROCESS

Option	Description
conf	Path and name of the authGateway configuration file. The default file is <b>\$RADIUSDIR/conf/authGateway.conf</b> .
name	Name of the authGateway process.

The MML options listed in [Table 45 on page 349](#) are still supported for backward compatibility.

**NOTE:** Starting from SBR Carrier 8.4.0 release, the **authGateway.conf** file configuration takes precedence over any existing MML CREATE-PROCESS options. While upgrading to SBR Carrier Release 8.4.0 or later, the **authGateway.conf** file will be imported from older versions of SBR Carrier. If there is no specified configuration present in the imported **authGateway.conf** file, SBR Carrier uses the existing MML configurations. If no MML configurations are present, the **authGateway.conf** defaults are used. If any mandatory parameters (port, host, and node) are missing, then error messages are logged.

Table 45: authGateway Process Options Supported for Backward-Compatibility

Option	Description
appctx	MAP protocol revision (2 or 3). Only MAPv3 retrieves quintets, so it must used to support EAP-AKA.
conf	Path and name of the authGateway configuration file. The default file is <b>\$RADIUSDIR/conf/authGateway.conf</b> .

Table 45: authGateway Process Options Supported for Backward-Compatibility (continued)

Option	Description
debug	Sets a debug level. Use the following:  <b>-debug 0xff</b>
host	Local hostname. Use the hostname associated with the IP address that the authGateway listen on, and ensure that the entry is coordinated with the <b>radius/GWrelay.conf</b> file. If a hostname is not specified, 0.0.0.0 is used.
invkretry	Number of invoke retry.
invktimeout	Duration of invoke timeout in seconds.
lgti	(Global Title only) Local GTI value, usually 4 for C7 and 2 for A7.
lmsisdn	(Global Title only) MSISDN of this local node.
lnai	(GT only) Nature of Address Indicator. Indicates the scope of the address value, such as whether it is an international number (includes country code) or a national number (no country code).  1 Subscriber Number—no area code (example: 5551234)  2 unused  3 National Significant Number—no country code (example: 2015551234)  4 International Number—includes country code (example: 12015551234)



Table 45: authGateway Process Options Supported for Backward-Compatibility (*continued*)

Option	Description
lnp	<p>(Global Title only) Local Numbering Plan.</p> <p>Acceptable values are:</p> <p>1—ISDN/Telephony</p> <p>3—DATA</p> <p>4—TELEX</p> <p>5—Maritime Mobile</p> <p>6—Land/Mobile</p> <p>7—ISDN/Mobile</p> <p>10—British Telecom special 1</p> <p>11—British Telecom special 2</p> <p>14—Private Network</p>
lpc	Local Point Code (PC).
lri	Routing indicator - 0 for GT (Global Title), 1 for PC/SSN.
lssn	Local Subsystem Number (SSN) (required).
ltn	(Global Title only) Local Translation Type. Generally in a live network TT is always 0.
max_requests	The maximum number of simultaneous MAP dialogs.
monitor	Activates Message Activity Monitor.
name	Name of the process.
no_rst	Disables automatic restart of process.
node	Node name.
port	Port number used by the SCTP association with the client.
prot	Variant used (C7, A7, or CH7).

Table 45: authGateway Process Options Supported for Backward-Compatibility (*continued*)

Option	Description
rssn	Subsystem number of HLR.
trace	We recommend setting this to 0xff; this enables debug tracing and displays the trace information about the console. (Consists of a trace of all MAP messages that are formatted and sent down the stack.)  Use the tracefile option to capture the trace information to a file.
tracefile	Captures the trace information to a file. The filename follows the <b>-tracefile</b> switch. Include the directory in the filename.

**Example—Creating and Starting the authGateway Process**

Notice that the SSN=7 in the **CREATE-CPC** of the previous example (Example 2 in [“Configuring the authGateway Routing Location Information” on page 338](#)) becomes the Issn (local subsystem number) in the **CREATE-PROCESS** command of this example. The SSN=6 in the **CREATE-REMSSN** command of the previous example (Example 2 in [“Configuring the authGateway Routing Location Information” on page 338](#)) becomes the -rssn (remote subsystem number) in this example. We recommend that you specify an absolute (full) path in the **EXEC** command.

The following configuration example explains how to create and start three authGateway instances:

(Lines are wrapped.)

```
CREATE-PROCESS:NAME="GMT", CE="sbr-lnx-perf", EXEC="/opt/JNPRsbr/radius/authGateway
-name GMT -conf /opt/JNPRsbr/radius/conf/authGateway.conf -port 2003"

START-PROCESS:NAME="GMT",CE="sbr-lnx-perf";

CREATE-PROCESS:NAME="GMT1", CE="sbr-lnx-perf", EXEC="/opt/JNPRsbr/radius/authGateway
-name GMT1 -conf /opt/JNPRsbr/radius/conf/authGateway.conf -port 2005"

START-PROCESS:NAME="GMT1",CE="sbr-lnx-perf";

CREATE-PROCESS:NAME="GMT2", CE="sbr-lnx-perf", EXEC="/opt/JNPRsbr/radius/authGateway
-name GMT2 -conf /opt/JNPRsbr/radius/conf/authGateway.conf -port 2007"

START-PROCESS:NAME="GMT2",CE="sbr-lnx-perf";
```

**NOTE:** MML commands are saved in MML files that can be loaded into Signalware. See [“Loading the MML Configuration Settings” on page 354](#).

## Configuring the GWrelay.conf File

The GWrelay application is used to pass authentication requests between SBR Carrier and the authGateway instances in a round-robin method. The **GWrelay.conf** file is used to define the source and destination ports through which an SCTP connection is established between the GWrelay application and authGateway instances.

You can modify the LOCAL\_HOST, REMOTE\_HOST, and RELAY\_SERVER lines in the **GWrelay.conf** file to define DNS names and port numbers. When you specify a DNS name for a local or remote host, you can enter the host's IP address in brackets as a backup. We recommend that you make hostname and IP address entries in the **/etc/hosts** file because it is more reliable than DNS.

### Example—Configuring the GWrelay.conf File

The following example explains how to define source and destination ports for three authGateway instances:

```
LOCAL_HOST sbr-blr-vm5:2002
REMOTE_HOST sbr-blr-vm5:2003 [10.20.0.2]

LOCAL_HOST sbr-blr-vm5:2004
REMOTE_HOST sbr-blr-vm5:2005 [10.20.0.2]

LOCAL_HOST sbr-blr-vm5:2006
REMOTE_HOST sbr-blr-vm5:2007 [10.20.0.2]

RELAY_SERVER sbr-blr-vm5:2000
```

**NOTE:** The specified host-name and port parameters in the REMOTE\_HOST line must match the -host and -port options in the MML CREATE-PROCESS statement, respectively.

## Starting the GWrelay Process

You can use the **sbrd** script to start and stop the GWrelay process. All **sbrd** commands can be executed only by the root user. To start the GWrelay process, execute **./sbrd start GWrelay**. To stop the GWrelay process, execute **./sbrd stop GWrelay**. To restart the GWrelay process, execute **./sbrd restart GWrelay**.

**NOTE:** The GWrelay application gets terminated automatically when all the configured authGateway instances are down. So, you must manually start the GWrelay application when the authGateway instances are restarted.

If you have set the **GWRELAYENABLE** parameter in the **sbrd.conf** file to 1 or answered **Yes** to the question **Do you want to enable "GWrelay" Process? [n]:** while running the SBR Carrier configuration script, then the GWrelay process will be started, stopped, or restarted when you execute the **./sbrd start**, **./sbrd stop**, or **./sbrd restart** script respectively.

## Configuring the ulcmmg.conf File

The **ulcmmg.conf** file establishes the connection between the GWrelay application and SBR Carrier.

The **ulcmmg.conf** file consists of two lines, as shown in the following example. Modify the **ulcmmg.conf** file shipped with SBR Carrier so that hostnames of LOCAL\_HOST and REMOTE\_HOST are same. If you specify a DNS name for a local or remote host, you can enter the host's IP address in brackets as a backup. Making an entry in the **/etc/hosts** file is recommended because it is more reliable than DNS.

### Example

```
LOCAL_HOST myhost.com:2001
REMOTE_HOST myhost.com:2000 [172.25.97.230]
```

**NOTE:** If an IP address is specified, it must be the address of the server specified as the HostName set when the Create-Process -host option is invoked.

For additional examples, see [“Sample authGateway Command and File” on page 355](#).

## Loading the MML Configuration Settings

The files containing MML commands need to be loaded into Signalware. Enter an SWMML command for each **.mml** file that you created.

In the following example, three files are loaded into Signalware.

- The **links.mml** file sets up the links, link sets, and routes.
- The **config\_authgateway.mml** file configures the authGateway application.
- The **start\_authgateway.mml** contains authGateway startup information.

These files do not exist and are used here as an example. You need to create your own files for accomplishing the tasks of creating links, configuring the authGateway application, and starting the authGateway application.

### Example

```
$ swmml -f links.mml
$ swmml -f config_authgateway.mml
$ swmml -f start_authgateway.mml
```

**NOTE:** Signalware retains the MML configuration commands and uses the same configuration each time you start Signalware. To reconfigure the nodes, use the **configureNodes** command with the **-f** or **-clean** or **-realclean** options.

## Sample authGateway Command and File

### sampleCreateProcess.mml

```
CRTE-PROCESS:NAME="GMT",CE="as1",EXEC="authGateway -name GMT -port 2000 -host as1
-node MGW -prot C7 -conf conf/authGateway.conf.100 -lri 1 -lssn 7 -appctx 3
-ldigits 358402114300"; START-PROCESS:NAME="GMT", CE="as1";
```

### authGateway.conf (Gateway Routing Configuration File)

Each line of this configuration file describes:

- The parameters used to form the destination SCTP address.
- The decimal value of the bearer service (bs) and/or teleservices (ts) indicating an authorization.
- A string associated with each bs or ts value. If no bs and no ts are specified, no authorization is performed by the authGateway.

For each request, the first digits of the IMSI are compared with **odigits**. The first line of the configuration file that matches is selected for the current request. If the routing indicator (rri) specifies that the routing is done on GT, the leading digits are replaced with the new digits (ndigits) to perform the numbering plan translation. If a parameter is not present on the line, it is not present in the destination SCTP address. If

this is incompatible with the routing indicator, an error message is issued when the MAP authentication Gateway is started.

After modifying this configuration file, the authGateway has to be restarted. Modification of the **ndigits** value may require the creation of a new GT rule (CREATE-GT command).

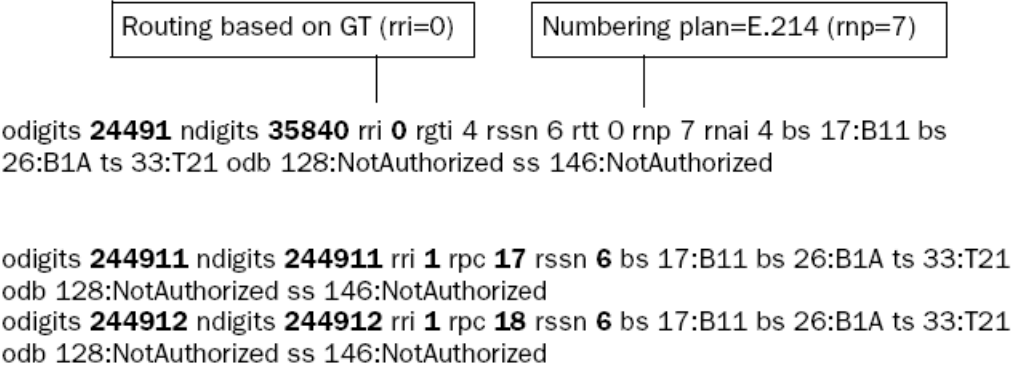


Table 46: Parameters Used in Create and Start Commands

Parameter	Description
rri	Routing indicator - 0 for GT (Global Title), 1 for PC/SSN (Point Code/Subsystem Number).
rgti	Local Global Title Indicator value. 4 for C7; 2 for A7. (Usually 4.)
rsn	Subsystem Number of HLR.
rtt	(GT only) Translation Type (usually 0).
rnp	(GT only) Numbering Plan. 1=E.164=ISDN-TEL, 7=E.214=ISDN-MOB.
rnai	Nature of address indicator. (4=INT)
bs	<p><b>bs dec:string</b></p> <p>If the specified bearer service exists in HLR SIM-profile, then <b>string</b> is returned for further processing (see ProfileMap in <b>simauth.aut</b>).</p>

Table 46: Parameters Used in Create and Start Commands (*continued*)

Parameter	Description
ts	<p><b>ts dec:string</b></p> <p>If the specified teleservice exists in HLR SIM-profile, then <b>string</b> is returned for further processing (see ProfileMap in simauth.aut).</p>

# 7

PART

## Uninstalling Steel-Belted Radius Carrier Software

---

[Removing Steel-Belted Radius Carrier Software](#) | 359

---



# Removing Steel-Belted Radius Carrier Software

## IN THIS CHAPTER

- [Uninstalling the Steel-Belted Radius Carrier Server Software | 359](#)

This chapter explains how to remove all Steel-Belted Radius Carrier Release 8.6.0 software from a server.



**CAUTION:** Do not uninstall Steel-Belted Radius Carrier if you intend to install another version of the Steel-Belted Radius Carrier software. If you remove an existing installation, you cannot migrate that data and configuration information forward to the newer version.

These topics are in the chapter:

## Uninstalling the Steel-Belted Radius Carrier Server Software

The software uninstall procedure is essentially the same for a standalone server or cluster node, the only difference being how the processes are stopped.

To uninstall the Steel-Belted Radius Carrier server software from its default location (`/opt/JNPRsbr`) on Solaris server:

1. Log in to the server or node as root.
2. Stop the RADIUS process. The exact syntax for this step depends on the server type:

For a standalone server, execute:

```
cd /opt/JNPRsbr/radius
./sbrd stop
```

In a cluster environment, stopping the SSR and RADIUS processes on the various node types, must be done in a specific order. Follow the procedures described in

[“When and How to Restart Session State Register Nodes, Hosts, and Clusters” on page 196](#), and then proceed to the next step.

3. (Optional) Back up your Steel-Belted Radius Carrier server directory.

Create a new archive directory to ensure that you do not overwrite an existing backup.

```
cd /opt/JNPRsbr
mkdir /opt/backups
tar cf - radius | ( cd /opt/backups; tar xfBp - )
```

4. Navigate to the directory where you unpacked the Steel-Belted Radius Carrier package, usually `/opt/JNPRsbr/radius/install`.

Execute:

```
cd /opt/JNPRsbr/radius/install/
```

5. Execute the **unconfigure** script:

Execute:

```
./unconfigure
```

6. A warning message is displayed. Enter **y** to proceed.

```
-----
SBR 8.60.50006 temporary cluster cambridge
on SunOS 5.10 Generic_141444-09 node sbrha-4(smdt)
is configured and processes are down, may be reconfigured if desired
-----
```

```
Unconfiguring Cluster Node...
```

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y
```

7. Clean up the user accounts.

A prompt is displayed to delete or preserve and clean up the `hadm` user account and associated directories.

- If you are going to reconfigure this server to serve as a standalone SBR Carrier server or become a node host in a cluster, enter **n**.
- If the server will be used for other purposes, enter **y**.

```

Locating OS user account and home directory...
hadm:x:16666:65536::/opt/JNPRhadm:/bin/bash
hadmg::65536:
drwxrwx---  2 hadm      hadmg      1536 Mar 28 10:59 /opt/JNPRhadm

```

```

WARNING: If you remove the OS user account hadm you will have
to recreate it, the associated OS group account hadmg, and
the associated home directory /opt/JNPRhadm
This is neither necessary nor recommended if you are updating
an existing configuration.  Remove the OS user account? [n]:

```

8. The script removes the software from various directories in /opt:

```

Cleaning directories
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRmysqld
/opt/JNPRndb_mgmd
/opt/JNPRndbd

```

9. Remove the software:

- a. Log in as root.
- b. Execute:
 

```

cd/opt
pkgrm JNPRsbr

```

Example:

```

root@life:/opt/JNPRsbr/radius/install> cd /opt
root@life:/opt> pkgrm JNPRsbr

```

The following package is currently installed:

```
JNPRsbr  JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Edition)
          (sparc) 8.60.50006
```

Do you want to remove this package? [y,n,?,q] **y**

```
## Removing installed package instance <JNPRsbr>
```

This package contains scripts which will be executed with super-user permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] **y**

```
## Verifying package <JNPRsbr> dependencies in global zone
```

```
## Processing package information.
```

```
## Executing preremove script.
```

Existing server directory will be backed up as:

```
/opt/JNPRsbr/radius/install/backups/2009:03:31-00:12:03
```

```
## Removing pathnames in class <none>
```

```
/opt/JNPRsbr/radius/install <non-empty directory not removed>
```

```
/opt/JNPRsbr/radius <non-empty directory not removed>
```

```
## Executing postremove script.
```

```
## Updating system information.
```

Removal of <JNPRsbr> was successful.

```
root@life:/opt> rm -rf /opt/JNPRsbr
```

To uninstall the Steel-Belted Radius Carrier server software from its default location (`/opt/JNPRsbr`) on Linux server:

1. Log in to the server or node as root.
2. Stop the RADIUS process. The exact syntax for this step depends on the server type:

For a standalone server, execute:

```
cd /opt/JNPRsbr/radius  
./sbrd stop
```

In a cluster environment, you need to stop the SSR and RADIUS processes on the various node types in a specific. Follow the procedures described in

[“When and How to Restart Session State Register Nodes, Hosts, and Clusters” on page 196](#), and then proceed to the next step.

3. (Optional) Back up your Steel-Belted Radius Carrier server directory.

Create a new archive directory to ensure that you do not overwrite an existing backup.

4. Navigate to the directory where you unpacked the Steel-Belted Radius Carrier package, usually **/opt/JNPRsbr/radius/install**.

Execute:

```
cd /opt/JNPRsbr/radius/install/
```

5. Execute the **unconfigure** script:

Execute:

```
./unconfigure
```

6. A warning message is displayed. Enter **y** to proceed.

```
-----  
SBR 8.60.50006 temporary cluster cambridge  
on Linux 3.10.0-693.el7.x86_64 #1 SMP Thu Jul 6 19:56:57 EDT 2017 node Cambridge  
is configured and processes are down, may be reconfigured if desired  
-----
```

```
Unconfiguring Cluster Node...
```

```
WARNING: You are about to make irreversible changes to this node.  
Are you sure that you wish to continue? (y,n): y
```

7. The script removes the software from various directories in /opt:

```

Cleaning directories
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRmysqld
/opt/JNPRndb_mgmd
/opt/JNPRndbd

```

## 8. Clean up the user accounts.

A prompt is displayed to delete or preserve and clean up the hadm user account and associated directories.

- If you are going to reconfigure this server so that it serves as a standalone SBR Carrier server or becomes a node host in a cluster, enter **n**.
- If the server will be used for other purposes, enter **y**.

```

Locating OS user account and home directory...
hadm:x:16666:65536::/opt/JNPRhadm:/bin/bash
hadmg::65536:
drwxrwx---  2 hadm      hadmg      1536 Mar 28 10:59 /opt/JNPRhadm

```

```

WARNING: If you remove the OS user account hadm you will have
to recreate it, the associated OS group account hadmg, and
the associated home directory /opt/JNPRhadm
This is neither necessary nor recommended if you are updating
an existing configuration.  Remove the OS user account? [n]:

```

```

Unconfigured

```

## 9. To remove the SBR software, execute:

- **yum remove sbr-cl** for the cluster configuration.
- **yum remove sbr-sa** for the standalone configuration.

## 10. The script removes the SBR software from the server.

# 8

PART

## Upgrading Your SSR Cluster

---

[Overview of Upgrading Your SSR Cluster | 366](#)

[Using a Transition Server to Mitigate Downtime While Upgrading Your Cluster | 369](#)

[Upgrading Your Cluster Using the Rolling Restart Method | 396](#)

[Upgrading Your Cluster Using the Backup, Destroy, and Re-Create Method | 447](#)

---

# Overview of Upgrading Your SSR Cluster

## IN THIS CHAPTER

- Upgrade Methods | 366
- Notes on Working with Data (D) Nodes | 367
- Using the SSR Configuration Script | 368

This section provides an overview of the procedures used to upgrade your SSR cluster. For information about upgrading standalone servers, see [“Installing and Configuring a SBR Carrier Standalone Server” on page 63](#).

These topics are included in this section:

## Upgrade Methods

SBR Carrier provides SSR cluster performance enhancements, as well as other product improvements. However, depending on the current release of software you are running, and which upgrade method you use to perform the upgrade, you may or may not be able to take advantage of the performance enhancements.

The choice of which upgrade method you use depends on your tolerance for downtime in the SSR cluster. You can use one of the following two methods to upgrade your SSR cluster:

- *Rolling restart upgrade*—Use this method to upgrade existing SBR Carrier release SSR clusters to Release 8.6.0. This upgrade provides numerous product improvements and corrections, but does not provide performance enhancements to the cluster. To experience product improvements and corrections, as well as improved performance in your cluster, you must upgrade using the backup, destroy, and re-create upgrade method. For complete details on the changes to the Release 8.6.0 software, see the *SBRC Release 8.6.0 Release Notes*. When using this upgrade method, we recommend you mitigate the risk of downtime in your cluster by using a transition server; for more information see [“Using a Transition Server to Mitigate Downtime While Upgrading Your Cluster” on page 369](#).
- *Backup, destroy, and re-create your cluster*—Use this method to upgrade from clusters running previous releases of SBR Carrier software to Release 8.6.0. This upgrade method results in the data nodes using



a multi-threaded process which improves performance over the single-threaded data node processes used in previous SBR Carrier software releases. In addition, Release 8.6.0 provides many improvements and corrections to the SBRC software. For complete details on the changes to Release 8.6.0 software, see the *SBRC Release 8.6.0 Release Notes*. Using the backup, destroy, and re-create upgrade method requires cluster downtime. For this reason, we recommend you use a transition server during the upgrade process. For details on using a transition server, see

[“Using a Transition Server to Mitigate Downtime While Upgrading Your Cluster” on page 369.](#)

## RELATED DOCUMENTATION

[Upgrading Your Cluster Using the Backup, Destroy, and Re-Create Method | 447](#)

[Upgrading Your Cluster Using the Rolling Restart Method | 396](#)

[Using a Transition Server to Mitigate Downtime While Upgrading Your Cluster | 369](#)

## Notes on Working with Data (D) Nodes

For performance reasons, the SSR ndbmt processes on data (D) nodes are configured to execute under the UNIX root account by default, as opposed to the UNIX hadm account. In particular, this allows the ndbmt processes to lock data in physical memory, which is faster, as opposed to allowing the OS to use swap space on disk, which is slower. The UNIX root account privilege is required in order to lock data in physical memory.

- The relevant configuration item is the **#sbrd-ndbd-run-as-root = true** parameter in the **[ndbd]** section of the **/opt/JNPRhadm/my.cnf** file. Note that the leading # character is required to distinguish this parameter as a **sbrd** script parameter; this parameter is *not* a comment and is always active. When the value of this parameter is true, the ndbmt processes execute under the UNIX root account. When the value of this parameter is false (or if the parameter is missing entirely), the ndbmt processes execute under the UNIX hadm account. The value of this parameter can only be changed immediately after configuring a data (D) node. The value of this parameter cannot be changed after the SSR processes are running.
- We recommend, although it is not necessary, that the parameter be configured the same on all data (D) nodes. In order to change the value of this parameter at a later time, you must unconfigure the data (D) node and then reconfigure it again.
- When the ndbmt processes are executed under the UNIX root account, it is extremely important that the **DataMemory** parameter in the **[ndbd default]** section of the **/opt/JNPRhadm/config.ini** file be configured properly with respect to the amount of physical memory that is actually available on the data (D) node. If the data (D) node does not have enough physical memory available, then the ndbmt processes can starve the entire machine, including the OS itself, for memory. By default, SBRC is configured under

the assumption that at least 8 GB of memory is available *solely* for ndbmysd processes. In practice, more than 8 GB is required to support the OS and other applications.

## Using the SSR Configuration Script

When running the SSR configuration script, option **2. Generate Cluster Definition**, always generates a 16-digit random number that is saved in **dbcluster.dat** as the CLUSTER\_DEFAULT\_KEY. Upon selecting configure script option **3. Configure Cluster Node**, the following occurs:

1. The CLUSTER\_DEFAULT\_KEY is used to configure **spi.ini** [Keys] if:
  - CurrentKey= has not already been configured.
  - and
  - The first key 1= has not already been configured.
2. Any IP addresses previously specified for nodes in the cluster are added to the **spi.ini** [Hosts] section if they are not already present.

**NOTE:** If you are migrating from an older **spi.ini**, then:

1. It is likely that [Keys] have already been configured and it is up to you to maintain [Keys] in this case.
2. Node IP addresses may still be added to [Hosts] but pre-existing (and possibly obsolete) addresses will not be removed.

Also note that the IP addresses specified for nodes during cluster configuration are primarily intended for SSR as opposed to RADIUS-specific configuration. If RADIUS traffic utilizes different IP addresses from SSR traffic, then it is up to the user to maintain **spi.ini** [Hosts].

# Using a Transition Server to Mitigate Downtime While Upgrading Your Cluster

## IN THIS CHAPTER

- [Cluster Migration Strategy | 369](#)
- [Cluster Migration Workflow | 372](#)
- [Creating a Transition Server | 373](#)
- [Installing the Cluster | 387](#)
- [Removing the Transition Server from Service | 390](#)
- [Cleaning the Transition Server | 390](#)

This chapter describes how to use a transition server to mitigate downtime while upgrading your cluster.

These topics are included in this section:

## Cluster Migration Strategy

The easiest way to replace an existing cluster with an upgraded cluster is to fully install and configure the new cluster and then simply cut over to the new install.

This causes a brief service disruption that can be mitigated if both clusters run online in parallel long enough for existing sessions to naturally drop off the old cluster as they end. Because no new sessions are added to the old cluster, after some period of time, most active sessions are managed by the new cluster. Any remaining long-term sessions are terminated when the old cluster is brought down. When they reconnect to the network, they connect to the new cluster.

Some sites may have a problem implementing this strategy because it requires enough servers to support two clusters, and not all sites have that many machines available.

## Using a Transition Server

To address this, we developed a migration strategy that uses a transition server. A single transition server temporarily takes the place of your entire existing cluster while the servers in your existing cluster are taken offline, upgraded with new SBRC Session State Register software, and then brought back online.

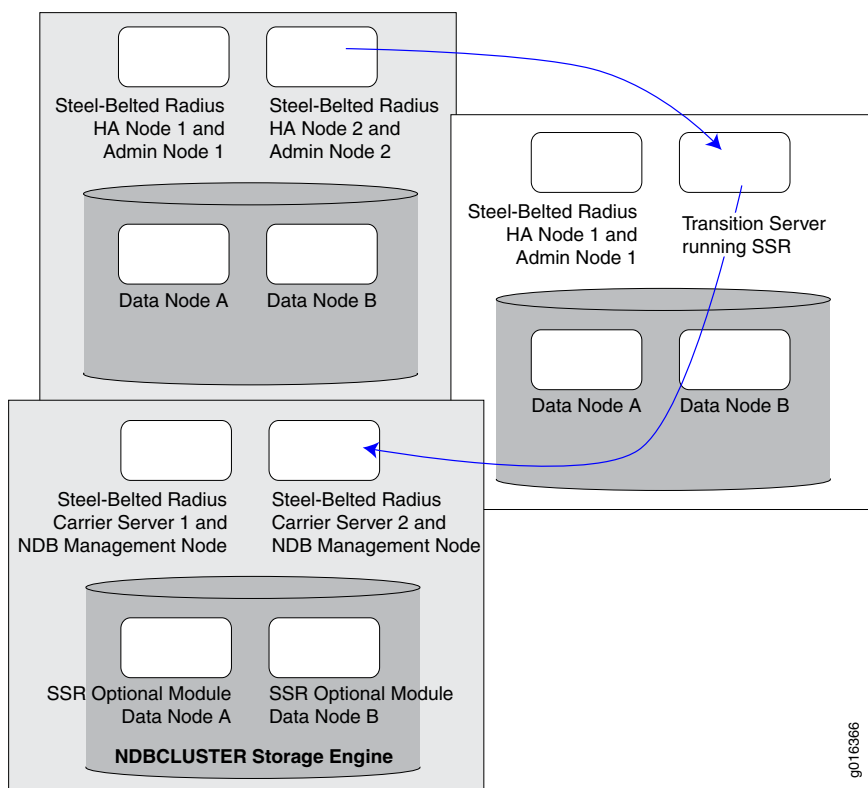
**NOTE:** We recommend using a transition server to mitigate downtime when performing any upgrade. The examples shown here depict the use of a transition server during an upgrade from SBR/HA 5.5 to a Release 8.6.0 Session State Register (SSR) cluster. However, we also recommend using a transition server when upgrading from previous releases of SBR Carrier software to Release 8.6.0.

The SBRC Temporary Cluster, also termed *the transition server*, is an exceptional node in the sense that it executes all processes on one machine. The transition server is assigned the node type= smdt. The **CreateDB.sh** and configuration of pool(s) must be done manually on a transition server, just like in a cluster.

Use a transition server (fifth machine) in addition to the four servers that a basic cluster installation requires to ensure redundancy. The additional fifth machine performs the work of the entire cluster while the four existing cluster servers are taken offline, updated, and brought online in a SSR Starter Kit configuration.

If an additional fifth machine is not available and you must work only with the four servers that currently make up your cluster, the transition server strategy can be adapted and one server in the existing cluster can be borrowed from the existing cluster and used as the transition server. This increases the risk of cluster failure during the switchover because some level of redundancy or capacity is removed from the existing, working cluster when a machine is taken offline.

Figure 18: SBR Migration Using the Transition Server Four-Server Strategy



#### Four-Server Strategy Only

If you must use the four-server strategy, look for special labels in the margin note area beside critical paragraphs in this chapter, like this one. They identify extra steps required if the transition server was part of an existing cluster and the cluster is running on three nodes. The step identified by the Four-Server Strategy Only label is not required if you use the five-node strategy (extra fifth machine) and leave the existing cluster fully functional.

The transition takes at least several hours and may take longer. The amount of time varies from site to site and depends upon:

- The number of servers involved.
- Whether the servers require a Solaris upgrade. Many SBR/HA Release 5.5 servers ran on Solaris 9 but Steel-Belted Radius Carrier Release 8.6.0 requires Solaris 11.3.36.10.0 or later.
- Whether a Release 8.6.0 test environment that replicates the production environment exists and can be moved to the production cluster, or the production environment needs to be created and tested.
- Whether one person or crew is installing one server at a time, or several are being installed in parallel.

## Individual Node Migration Guidelines

Before beginning an upgrade that reuses your existing cluster servers, confirm that all machines can meet the Release 8.6.0 server requirements in [“Before You Install Software” on page 39](#). In addition, observe these guidelines:

- Migration of SBR/HA Release 5.5 data nodes or management nodes is not supported. These nodes must be taken offline, existing software removed, and new software installed.
- No server may host both SBR/HA Release 5.5 software and Release 8.6.0 SSR software at the same time. Just one version may be installed at any time.
- Migration of SBR/HA Release 5.5 RADIUS node configuration files to Release 8.6.0 SBR Carrier (S) nodes is not supported because of the significant differences between SBR/HA Release 5.5 and Release 8.6.0.
- If an existing Release 8.6.0 environment exists, perhaps one used for testing, you can use the configuration files from those SBR Carrier (S) nodes for the production installation.

## Cluster Migration Workflow

The migration workflow follows a list of tasks that must be performed in a specific order to make the Session State Register cluster server fully functional. This guide presents groups of related tasks in sequential order.

The high-level view of the migration has several groups of tasks:

1. Planning—Because a SSR cluster involves four or more servers, planning the cluster topology and each server’s role and setting is essential. Planning is discussed, and a worksheet you can use to plan your installation is in [“Planning Your Session State Register Cluster” on page 32](#).
2. Creating a transition server—Setting up a transition server to temporarily provide cluster services while other servers are set up is discussed in this chapter. See [“Creating a Transition Server” on page 373](#).
3. Taking your existing cluster offline—After the transition server is handling traffic, the other nodes may be taken offline.
4. Preparing—The tasks to complete and the information to gather on each server in a cluster before software can be installed are discussed in detail in [“Before You Install Software” on page 39](#), and in the succeeding chapters that document the optional installation of Signalware that is required to support some features.

Because this chapter was written for new installations rather than migration, some tasks may need to be completed out of order, but all need to be completed.

5. Installing software—This includes completing preparatory steps that you were unable to complete while your existing servers were online, unpacking the SSR software package on all nodes, selecting

configuration tasks, and providing settings used to set up the cluster servers. The preparatory steps are discussed in [“Before You Install Software” on page 39](#); the actual software installation in [“Installing Session State Register Nodes” on page 94](#).

6. Configuring—Setting up configuration files for the SSR database is discussed in [“Customizing the SSR Database Current Sessions Table” on page 144](#).
7. [“Removing the Transition Server from Service” on page 390](#)—Taking the transition server offline and completing the cutover is discussed in this chapter. See [“Removing the Transition Server from Service” on page 390](#).

## Creating a Transition Server

To set up a transition server to temporarily take the place of your existing cluster, you need to prepare the server, install software, and configure the database.

### Preparing the Transition Server

To prepare the transition server:

1. Select the server.
  - The server must meet all the Release 8.6.0 hardware and software requirements listed in [“Before You Install Software” on page 39](#).
  - If the server is part of your existing cluster:
    - We recommend using the most powerful (the most RAM and greatest number of processors) server available because it processes a heavier-than-normal load during the transition.
    - We recommend using an SBR or management node, rather than a data node, to reduce front end processing on the old cluster and to maintain data redundancy.
  - If the server is acting as the transition server to be reconfigured as part of the new SSR Starter Kit cluster when it is reconfigured, it must be a combined SBR Carrier/management node host in a four-server cluster.

If you use Centralized Configuration Management to replicate SBR Carrier node configurations among a group of like nodes, the transition server cannot assume the role of primary CCM server in the new cluster because it is not the first SBR Carrier node to be configured.

### Four-Server Strategy Only

2. If the transition server is one of the existing cluster’s SBR (S) or management (M) nodes:

- a. Navigate to the **radius/install** subdirectory of the directory where the server package was installed.  
As root, stop the RADIUS processes:

Execute:

```
sbrd stop radius
```

- b. As root, identify and kill the **ndb\_mgmd** and **mysql** processes.

Execute:

```
Ps -ef|grep ndb_mgmd
```

```
kill -9 <Ndb_mgmd process_id>
```

```
Ps -ef|grep mysql
```

```
kill -9 <mysql process_id>
```

- c. Remove the Admin node and SBR node from the server.

Execute:

```
InstallAdminNode.sh - u
```

```
pkgrm JNPRsbr
```

- d. Perform all other tasks required to make the server conform to all installation prerequisites listed in [“Before You Install Software” on page 39](#).

## Unpacking and Configuring the New Software on the Transition Server

Before starting this procedure, review [“Before You Install Software” on page 39](#). In particular, review requirements for [“Setting Up External Database Connectivity \(Optional\)” on page 53](#) and [“Installing the SIGTRAN Interface \(Optional\)” on page 55](#). Steps in this procedure require the server to be preconfigured for these capabilities.

To unpack and configure the software on the transition server:

1. Log in as root.
2. Download, unpack, and install (**pkgadd -d**) the Steel-Belted Radius Carrier software package.  
See [“Unpacking Session State Register Software” on page 95](#) for complete instructions.
3. Navigate to the **radius/install** subdirectory of the directory in which the **JNPRsbr** package was installed (**/opt/JNPRsbr/radius/install** by default):

Execute:

```
cd /opt/JNPRsbr/radius/install/
```

4. Execute the **configure** script to set up the Steel-Belted Radius Carrier server software:

Execute:

```
./configure
```

5. Review and accept the Steel-Belted Radius Carrier license agreement.



Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

6. From the menu of configuration tasks, enter **5** to specify the type of installation as the **Create Temporary Cluster**.

```
Configuring SBR Software
```

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-4
is not configured and processes are down, needs to be configured
-----
```

```
1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.
```

```
2.  Generate Cluster Definition
    Creates new or updates existing cluster definitions.
    Modifies the shared directory but does not modify this node.
```

```
3.  Configure Cluster Node
    To be preceded by 'Generate Cluster Definition' on any node.
    Must be invoked on each and every node of the cluster.
```

```
4.  Reconfigure RADIUS Server
    Only on SBR nodes, updates the existing SBR configuration.
```

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q): **5**

#### 7. Specify the name of the cluster.

Enter the name exactly as you specified it in [Table 9 on page 35](#).

**Enter SBR cluster name [MyCluster]: MyCluster**

A warning prompt is displayed that explains the terms and limitations of the transition node.

#### 8. Enter the SSR Starter Kit license number, the license number for one SBR node, and a blank line.

While migrating to the new cluster, you are permitted to use the same licenses for the transition server as for the new cluster.

```
Enter Starter Kit license: XXXX XXXX XXXX XXXX
Enter SBR licenses meant only for this particular SBR node.
Enter one license per line and an empty line when finished.
Enter SBR full license: XXXX XXXX XXXX XXXX
Enter SBR feature license:
```

#### 9. Enter passwords for two internal accounts. The password input is not echoed to the screen; the fields appear to be blank.

```
All cluster nodes will share the same Session State Register (SSR).
Setting password for SSR admin account hadmsql
```

```

Password:
Again:
Setting password for SSR software account hadmsbr
Password:
Again:

```

10. The system generates the required configuration files and prompts you to view, accept, or reject them. Enter **a** to accept them and continue or **v** to view them.



**CAUTION:** We recommend that you enter an **r** to reject them only if a serious error was made when you provided information. We recommend that you not edit these files.

```

Reviewing Configuration Files...
-----

```

```

/opt/JNPRsbr/radius/install/tmp/dbcluster.rc
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen

```

```

View (v), accept (a), or reject (r) configuration files: a

```

11. Specify whether you want to use the JRE installed in your system to enable JDBC plug-ins and JavaScript implementation.

```

Do you want to configure Java Runtime Environment for JDBC Feature [n] :

```

- If no, press Enter to proceed to the next prompt. SBR Carrier does not support JDBC plug-ins unless you specify a valid JRE path.
- If yes, type **y** and press Enter. You are prompted to specify the path where the JRE is installed in your system. The Java Virtual Machine (JVM) architecture should be compatible with SBR Carrier.

**NOTE:** Java 1.8.0 or a later version is required to access the Web GUI. To support both JDBC plug-ins and Web GUI, it is recommended to use Java 1.8.0 or a later version with the JVM architecture compatible with your SBR Carrier. For example, if you are using the 64-bit version of SBR Carrier, you must use the 64-bit version of Java 1.8.0 or later.

```
Enter 64-bit libjvm.so path (Ex: /opt/jvm/jre/lib/amd64/server/ ) :
```

**NOTE:** If you enter an incorrect JVM path three times, SBR Carrier proceeds to the next step. In this case, you will not be able to use JDBC plug-ins. To specify the valid JVM path, you need to run the configure script again.

12. Supply the name of the initial admin user, root.

```
Enter initial admin user (UNIX account must have a valid password) [root]:
```

Press Enter to accept the default, root.

13. Do not set up centralized configuration management (CCM).

```
Enable Centralized Configuration Management (CCM) for this SBR node? [n]:
```

14. Specify whether you want to use the auto-restart module that automatically restarts the SBR Carrier server in case of an unexpected shutdown.

```
Do you want to enable "Radius WatchDog" Process? [n]: Y
Radius WatchDog feature set to Enable
Please ensure that Perl 5 or better is installed.
```

**NOTE:** If Perl version 5 is not installed, the radiusd script will not run, even if enabled by configuration, and SBR Carrier will operate without the auto-restart module running.

15. Specify whether you want to configure SBR Carrier to provide LDAP server emulation for configuration and statistics using the LCI.

```
Do you want to enable LCI? [n] :
```

- If no, press Enter to accept the default.
- If yes, enter y and press Enter. You are prompted to provide information for LCI configuration.
  - a. When you are prompted for the port number, enter the port number that is used for communication between SBR Carrier and the LDAP client.

**NOTE:** SBR Carrier uses port 667 as the default for LDAP emulation to avoid conflict with other LDAP servers.

```
Configure LCI Port [667]: 1026
```

- b. The script displays the interfaces available in the system. When you are prompted to enter interface addresses on which Steel-Belted Radius Carrier should listen for LCI requests, enter the addresses you want to use from the Available Interfaces list.

```
LCI Interface Configuration :
Available interfaces :
127.0.0.1
10.212.10.66

HELP : Enter one interface per line and an empty line when finished.

Enter LCI interface addresses from the above list.

Enter LCI interface address : 10.212.10.66
Enter LCI interface address : 127.0.0.1
Enter LCI interface address :
```

**NOTE:** SBR Carrier uses all interfaces for listening to LCI requests if you do not enter any interfaces.

- c. Specify whether you want to change the default LCI password to prevent unauthorized LDAP clients from accessing your database.

```
Do you want to change LCI Password? [n]:
```

- If no, press Enter to accept the default password.
- If yes, enter **y** and press Enter. You are prompted to enter a new password.

```
Do you want to change LCI Password? [n]: Y
```

```
Password must meet the following requirements:
```

1. 6-8 Alphanumeric characters.
2. No Special characters other than underscore ('\_').

```
Enter Password:
```

```
Confirm Password:
```

```
Password will be changed when SBR restarts.
```

**NOTE:** Make sure that the entered password is at least 6 alphanumeric characters and not more than 8 characters in length. The password should not include any special characters other than underscore ('\_').

**NOTE:** The configure script also checks whether the LDAP utilities (such as **ldapdelete**, **ldapmodify**, and **ldapsearch**) are installed in your system. For Linux, a warning message is displayed if you have not installed any of these utilities in your system. For Solaris, LDAP utilities are shipped with SBR Carrier package.

16. Specify whether you want to configure Steel-Belted Radius Carrier for use with an Oracle database.

To support this option, the server must already be configured as an Oracle client. (See [“Setting Up External Database Connectivity \(Optional\)” on page 53.](#))

```
Configuring for use with generic database
```

```
Do you want to configure for use with Oracle? [n]:
```

If no, press Enter to accept the default.

If yes, type **y** and press Enter. You are prompted for version and path information for the Oracle library files.

```

Do you want to configure for use with Oracle? [n]: y
Supported Oracle versions: 10, 11, 12
What version of Oracle will be used? [10]: 10
Configuring for use with Oracle 10
Setting the environment variable ORACLE_HOME
Enter ORACLE_HOME [/dbms/u10/app/oracle/product/10.2.0]:
Setting the environment variable LD_LIBRARY_PATH
Enter path for Oracle shared libraries [/dbms/u10/app/oracle/product/10.2.0/lib]:
Setting the environment variable TNS_ADMIN
Enter TNS_ADMIN [/dbms/u10/app/oracle/product/10.2.0/network/admin]:

```

**NOTE:** You must configure 64-bit Oracle client for 64-bit SBR Carrier.

17. Specify whether you want the Steel-Belted Radius Carrier server to communicate with an SS7 system using SIGTRAN.

To support this option, the server must already be configured to support SIGTRAN using Signalware. (See [“Installing the SIGTRAN Interface \(Optional\)” on page 55](#) for an overview, and [“SIGTRAN Support for Steel-Belted Radius Carrier” on page 309](#) for specific instructions.)

```

Do you want to configure for use with SIGTRAN? [n]: y
Configuring for use with SIGTRAN
Setting the environment variable OMNI_HOME
Enter OMNI_HOME [/opt/JNPRss7]:

```

18. Specify whether you want to start the GWrelay process while executing the `./sbrd start` script.

```

Do you want to enable "GWrelay" Process? [n]: y
GWrelay will be started with sbrd

```

19. Specify whether you want to install the optional SNMP module to monitor your Steel-Belted Radius Carrier server from an SNMP management station.

**Do you want to configure SNMP? [n]:**

- If no, press Enter to proceed to the next prompt.
- If yes, type **y** and press Enter. The installer prompts you for the information it needs to configure the `jnbrsnmpd.conf` and `startsnmp.sh` files.

- a. When you are prompted for a community string, enter the community string used to validate information sent from the SNMP subagent on the Steel-Belted Radius Carrier server to your SNMP management station.

**Choose a community string: public**

- b. When you are prompted for a range of IPv4 addresses, specify a starting IP address in Classless Inter-Domain Routing format. To specify that only one host may query the agent, enter the IP address of the host followed by /32. To specify that any host on a designated class C network may query the agent, enter the starting address of the network followed by /24.

**Specify the range of IPv4 addresses that may query this agent, such as 1.2.3.0/24.**

**Address range: 192.168.70.0/24**

- c. If you are using SNMPv2, enter the DNS name or IP address of the trap sink to receive trap information from the SNMP subagent on the Steel-Belted Radius Carrier server.

**SNMPv2 trap sink: 192.168.70.86**

- d. Set the SNMP agent port.

Although you may specify the default SNMP port, 161, we recommend that you specify a different port to avoid contention with other agents that are likely to already be using 161. If you choose an alternate port, make a note of it because your MIB browser needs to be configured to the same setting.

**Specify SNMP agent listening port[161]: 24161**

- e. Specify a trap sink address, if required.

**Optionally specify a trap sink address that will receive SNMPv2 trap**

**[localhost]: 172.28.72.83 2**

**SNMPv2 trap sink port[162]:**

**Configuration of SNMP complete.**

20. The script searches for the Java 1.8.0 or later version in the default system path and displays a confirmation message if found.

```
Configuring Admin GUI Webserver
Compatible Java version 1.8.0_66 found in: /usr/java/jdk1.8.0_66
```

If the specific version is not found, the script prompts you to enter the path where the specific Java version is installed in your system.

```
Enter Java version 1.8 installed path :
```

21. Specify whether you want to install a custom SSL certificate for the Web GUI.



```
Do you want to install custom SSL certificate for Admin WebServer? [n]:
```

- If no, press Enter. A self-signed certificate is created and installed in your server.
- If yes, enter **y** and press Enter. You are prompted to enter the absolute path where the SSL certificate is available. For example, **/opt/customSSLCert.pfx**.

```
Enter the absolute path to certificate.  
Note: Only *.pfx files are accepted. (Example-/opt/customSSLCert.pfx):
```

When you are prompted for the password, enter the password to open the SSL certificate.

```
Enter the password to open the certificate :
```

22. Specify whether you want to configure the Steel-Belted Radius Carrier server to autoboot (restart automatically when the operating system is restarted). We recommend that you enable autoboot behavior.

**Enable (e), disable (d), or preserve (p) autoboot scripts [e]: e**

A local **/radiusdir/radius/sbrd** script is always created, and **/opt/JNPRhadm/sbrd** is always a symbolic link to this local copy.

- If you enter **e** (enable), the **configure** script copies the local **sbrd** script to **/etc/init.d**, where it is automatically invoked by the OS whenever the OS is stopped or started.
- If you enter **d** (disable), the **configure** script removes all copies of the **sbrd** script from **/etc/init.d**, thus, disabling autoboot for all versions of Steel-Belted Radius Carrier.
- If you enter **p** (preserve), the **configure** script nothing, thereby leaving your previous autoboot scripts unchanged.

When you finish entering settings, the script configures Steel-Belted Radius Carrier with the specified settings and then displays:

```
The SBR Admin Web GUI can be launched using the following URL:  
https://<servername>:2909  
  
Configuration complete
```

23. Enter **q** to end the script.

```
-----
SBR 8.60.50006 temporary cluster cambridge
on SunOS 5.10 Generic_141444-09 node sbrha-4(smdt)
is configured and processes are down, may be reconfigured if desired
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

#### 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on any node.

Must be invoked on each and every node of the cluster.

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

```
READY: last operation succeeded, created temporary cluster.
Enter the number of the desired configuration task or quit (4,q): q
```

## Configuring the Transition Server

After the new software is configured, you need to configure the transition server as a temporary replacement for your existing cluster. All cluster traffic is ultimately switched to this single transition server temporarily, while you take the other nodes in the existing cluster down and upgrade and reconfigure them. So, you need to configure the temporary transition server as close to your existing configuration as possible. To do this, you need to configure the server configuration files on the temporary transition server for your environment.

We recommend that you complete the configuration of all server initialization (.ini) files, authentication (.aut) files, accounting (.acc) files, as well as configure any proxy setup you may require. Carefully review the *SBR Carrier Reference Guide* and configure all files for your environment prior to starting the temporary transition server.

Also review the *SBR Carrier Administration and Configuration Guide*, and plan the configuration steps for your particular environment. You cannot connect to the transition server with Web GUI until the RADIUS process is started; however, we recommend you plan out the administration of the server before starting the RADIUS process.

After you have completed the configuration of the various configuration files described in the *SBR Carrier Reference Guide*, remember to come back and complete this procedure.

1. Start the **ssr** process.

As root, execute:

```
cd /opt/JNPRsbr/radius
./sbrd start ssr
```

Status messages are displayed as the programs start:

```
Starting ssr management processes
Starting ssr auxiliary processes
Starting ssr data processes
```

2. Create the session database.

Log on as hadm and execute:

```
./CreateDB.sh
```

If you need to customize the sessions database to match your existing session database, see [“Customizing the SSR Database Current Sessions Table” on page 144](#).

3. (Optional) If you are using IP address pools in your existing cluster, you may need to either configure new IP address ranges, or taper off the use of a range from your existing cluster by removing a range and waiting for the addresses to be released to avoid giving out conflicting IP addresses to multiple users. For details on setting up IP address pools and ranges using the SSR Administration Scripts, see the *SBR Carrier Administration and Configuration Guide*.

**NOTE:** We recommend you check with the Juniper Networks Technical Assistance Center (JTAC) if you are using IP address pools and setting up a transition server.

4. Navigate to the **radius/install** subdirectory of the directory in which the **JNPRsbr** package was installed (**/opt/JNPRsbr/radius/install** by default), and start the RADIUS process:

As root, execute:

**sbrd start radius**

Status messages are displayed as the programs start:

```
Starting radius server processes
RADIUS: Process ID of daemon is 13224
RADIUS: Starting DCF system
RADIUS: Configuration checksum: 2D D6 38 1D
RADIUS started
.
.
.
RADIUS: DCF system started
```

5. Finish configuring the transition server using Web GUI. Follow the steps outlined in [“Basic SBR Carrier Node Configuration” on page 138](#). For complete details, see the *SBR Carrier Administration and Configuration Guide*.

After you have finished configuring the temporary transition server and you are sure it is configured properly to handle all traffic, you can switch all traffic to the transition server. See [“Switching Traffic to the Transition Server” on page 387](#).

## Switching Traffic to the Transition Server

After the transition server is set up and tested, and a working database created, reconfigure the site's routers to gradually direct traffic to the transition server instead of to the existing cluster's SBR servers.

When the transition server is handling all traffic, back up your existing cluster servers and take them offline.

## Installing the Cluster

With your existing cluster servers offline, you can proceed with a normal cluster installation:

1. Confirm that the cluster topology plan is complete. See [“Planning Your Session State Register Cluster” on page 32](#).

### Four-Server Strategy Only

If you are using the four-server strategy and plan to incorporate the transition server into the new Starter Kit cluster, the transition server must be the second SBR/management (SM) node host.

**NOTE:** During this procedure, the new cluster is configured as if the second SBR/management node host were present. However, it is not present because it is currently functioning as the transition server, so you defer configuring and starting the second SBR/management node host until after the new cluster is operational.

2. Make sure that each server in the new cluster conforms to the requirements in [“Before You Install Software” on page 39](#).



**CAUTION:** Do not skip this step; the server requirements for Session State Register have changed significantly since SBR/HA Release 5.x.

3. Install and configure the software on all cluster nodes.

Follow the procedures in [“Installing Session State Register Nodes” on page 94](#).

### Four-Server Strategy Only

If you are using the four-server strategy and plan to incorporate the transition node as the second SBR/management node in the new cluster, skip [“Setting Up the Second SBR/Management Node in a Starter Kit” on page 114](#) entirely as you work through the node installations.

4. When you begin configuring the cluster nodes, if you edited the **CurrentSessions.sql** file on the transition server during the procedure for [“Configuring the Transition Server” on page 385](#), you can copy that

**CurrentSessions.sql** to the first management node that you set up. See [“Customizing the SSR Database Current Sessions Table” on page 144.](#)

5. Start the cluster.

- If you are using a five-node strategy, use the [“Initial Cluster Startup Procedure” on page 132.](#)
- In the following procedure, each time the **sbrd status** command is executed results similar to this example are displayed:

```
hadmUser$>/opt/JNPRsbr/radius/sbrd status
```

```
[ndbd(NDB)]      2 node(s)
id=10    @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11    @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1      @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2      @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
```

```
[mysqld(API)]   4 node(s)
id=21     @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=22     @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=30     @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=31     @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
```

Examine the line starting with **id=**, and verify that there are no references to **starting**, **connecting**, or **not connected**. Any of these references indicate the process has either not finished starting, or the node is not connected properly. You may need to execute the **sbrd status** command more than once because it only shows a snapshot of activity; the display does not refresh automatically. Do not proceed to the next node until you are sure the process has started properly and the node is connected.

#### Four-Server Strategy Only

- If you are using the four-server strategy, start the new cluster using a non-standard series of commands because the fourth server that hosts the second management node is missing. Use this sequence of commands:

- a. Log in to the *SBR/management node* as root.
  - b. Change directories to **/opt/JNPRsbr/radius/**.
  - c. Execute:  
**./sbrd start ssr**
  - d. Execute:  
**./sbrd status**
  - e. Examine each line and ensure the SSR process is running without error.
  - f. Log in to a *data node* as root.
  - g. Change directories to **/opt/JNPRsbr/radius/**.
  - h. Execute:  
**./sbrd start ssr**
  - i. Execute:  
**./sbrd status**
  - j. Examine each line and ensure the SSR process is running without error.
  - k. Log in to the second *data node* as root.
  - l. Change directories to **/opt/JNPRsbr/radius/**.
  - m. Execute:  
**./sbrd start ssr**
  - n. Execute:  
**./sbrd status**
  - o. Examine each line and ensure the SSR process is running without error.
  - p. Go back to the *management node*, still as root.
  - q. Change directories to **/opt/JNPRhadm/**.
  - r. Log in as **hadm**.
  - s. Execute:  
**./CreateDB.sh**
6. Run **CreateDB.sh** on each SBR/management node and each management node in the cluster.
- If you need to customize the sessions database, see  
[“Customizing the SSR Database Current Sessions Table” on page 144.](#)
7. Configure at least one IP address pool and one range using the SSR Administration Scripts. See  
[“Testing the Installation with DemoSetup.sh” on page 482.](#) See the *SBR Carrier Administration and Configuration Guide* for details on configuring IP address pools and ranges.

**NOTE:** We recommend you consult with the Juniper Networks Technical Assistance Center (JTAC) if you are using IP address pools and creating a transition server.

8. Start the RADIUS process on the SBR/management node.
  - a. Log in as root to the SBR/management (sm) node.
  - b. Change directories to `/opt/JNPRsbr/radius/`.
  - c. Execute:  
`./sbrd start radius`
  - d. Execute:  
`./sbrd status`
  - e. Examine each line and ensure the RADIUS process is running without error.

Now that the RADIUS process is running, you can complete the configuration using Web GUI. See [“Basic SBR Carrier Node Configuration” on page 138](#). For complete details, see the *SBR Carrier Administration and Configuration Guide*.

## Removing the Transition Server from Service

After you bring the new cluster online, configure it, and test it, begin transferring live traffic to it and away from the transition server. When all traffic has been shifted to the new cluster and the number of ongoing sessions managed by the transition server has reached a suitably low level, take the transition server offline. Some sessions are terminated, but reconnect through the new cluster.

## Cleaning the Transition Server

To free the licenses used by the transition server, and clean up installed software, uninstall the SBR Carrier software. See [“Uninstalling Steel-Belted Radius Carrier Software” on page 358](#).

1. Log in as root.
2. Navigate to the directory where you installed the Steel-Belted Radius Carrier package, usually `/opt/JNPRsbr`. Then, navigate to the `/opt/JNPRsbr/radius` subdirectory.
3. Stop the RADIUS processes.  
  
Execute:  
`./sbrd stop radius`



4. Stop the **ssr** processes:  
`./sbrd stop ssr`
5. Navigate to the `/opt/JNPRsbr/radius/install` directory and run the **configure** script:  
 Execute:  
`./configure`
6. From the menu of configuration tasks, enter **1** to specify **Unconfigure Cluster Node**.

Configuring SBR Software

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-4
is not configured and processes are down, needs to be configured
-----
```

```
1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.
```

```
2.  Generate Cluster Definition
    Creates new or updates existing cluster definitions.
    Modifies the shared directory but does not modify this node.
```

```
3.  Configure Cluster Node
    To be preceded by 'Generate Cluster Definition' on any node.
    Must be invoked on each and every node of the cluster.
```

```
4.  Reconfigure RADIUS Server
    Only on SBR nodes, updates the existing SBR configuration.
```

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q): **1**

#### 7. A warning message is displayed. Enter **y** to proceed.

```
-----
SBR 8.60.50006 temporary cluster cambridge
on SunOS 5.10 Generic_141444-09 node sbrha-4(smdt)
is configured and processes are down, may be reconfigured if desired
-----
```

Unconfiguring Cluster Node...

WARNING: You are about to make irreversible changes to this node.  
Are you sure that you wish to continue? (y,n): **y**

#### 8. Clean up the user accounts.

A prompt is displayed to delete or preserve and clean up the hadm user account and associated directories.

- If you are going to reconfigure this server to serve as a standalone SBR Carrier server or become a node host in a cluster, enter **n**.

- If you want the server to be used for other purposes, enter **y**.

```

Locating OS user account and home directory...
hadm:x:16666:65536::/opt/JNPRhadm:/bin/bash
hadmg::65536:
drwxrwx---  2 hadm      hadmg      1536 Mar 28 10:59 /opt/JNPRhadm

```

```

WARNING: If you remove the OS user account hadm you will have
to recreate it, the associated OS group account hadmg, and
the associated home directory /opt/JNPRhadm
This is neither necessary nor recommended if you are updating
an existing configuration.  Remove the OS user account? [n]:

```

9. Enter **q** to end the script.

```

Cleaning directories
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRmysqld
/opt/JNPRndb_mgmd
/opt/JNPRndbd

```

```

-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-4
is not configured and processes are down, needs to be configured
-----

```

```

1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.

```

## 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.

## 3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on any node.  
Must be invoked on each and every node of the cluster.

## 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

## 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

## 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

READY: last operation succeeded, node unconfigured.  
Enter the number of the desired configuration task or quit (2,q): q

### Four-Server Strategy Only

10. You can now set up the server as the second SBR/management node in the new cluster. Because all preparatory steps have already been performed, follow the steps in [“Setting Up the Second SBR/Management Node in a Starter Kit” on page 114](#), with the exception of the last step for configuring the data nodes. After completing these steps, remember to return to this procedure to complete the remaining steps.
11. Start the second SBR/management node using this sequence of commands:
  - a. Log in to the SBR/management node as root.
  - b. Change directories to `/opt/JNPRsbr/radius/`.
  - c. Execute:  
`./sbrd start ssr`
  - d. Execute:  
`./sbrd status`
  - e. Examine each line and ensure the SSR process is running without error.
12. Create the database on the SBR/management (sm) node:
  - a. Log in as **hadm**.
  - b. Execute:  
`./CreateDB.sh`  
  
If you need to customize the sessions database, see [“Customizing the SSR Database Current Sessions Table” on page 144](#).
13. Start the RADIUS process on the SBR/management node:
  - a. Log in as root.
  - b. Change directories to `/opt/JNPRsbr/radius/`.
  - c. Execute:  
`./sbrd start radius`
  - d. Execute:  
`./sbrd status`
  - e. Examine each line and ensure the RADIUS process is running without error.
14. Now that the RADIUS process is running, you can complete the configuration using Web GUI. See [“Basic SBR Carrier Node Configuration” on page 138](#). For complete details, see the *SBR Carrier Administration and Configuration Guide*.

# Upgrading Your Cluster Using the Rolling Restart Method

## IN THIS CHAPTER

- Overview of the Rolling Restart Upgrade Method | 396
- Preparation | 399
- Upgrading the M Nodes in the Cluster | 400
- Upgrading the SM Nodes in the Cluster | 409
- Upgrading the S Nodes in the Cluster | 422
- Upgrading the Data (D) Nodes with the New Software | 431
- Launching Web GUI | 446

This section describes the rolling restart upgrade procedure, which you can use to upgrade servers running in a SSR cluster. For information about upgrading servers running standalone, see [“Installing and Configuring a SBR Carrier Standalone Server”](#) on page 63.

These topics are included in this section:

## Overview of the Rolling Restart Upgrade Method

This section describes how to perform a rolling restart upgrade to your 8.6.0 SSR cluster. By following this procedure, you should not experience any downtime in your cluster.

**NOTE:** Although you should not experience any service interruption during this procedure, we recommend you schedule a maintenance window when performing this procedure. To mitigate the risk of experiencing downtime in your cluster, we recommend using a transition server. For instructions on creating and using a transition server, see [“Using a Transition Server to Mitigate Downtime While Upgrading Your Cluster”](#) on page 369.

**NOTE:** The example procedures in this section assume your cluster consists of one SM node, one S node, two M nodes, and two D nodes, as follows:

- First M node, bng-mars.englab.juniper.net (10.212.10.68)
- Second M node, bng-sbr-perf1 (10.212.10.66)
- SM node, sbr1.englab.juniper.net (10.212.10.213)
- First D node, bng-sbr-perfm3000-3 (10.212.10.188)
- Second D node, bng-sbr-perf2 (10.212.10.67)
- S node, bng-sbrha-1 (10.212.10.65)

## Summary of the Rolling Restart Upgrade Method

During this upgrade procedure, you will need to:

1. Stop the SSR management process on the M node that you are planning to upgrade, upgrade the M node (install and configure the new software), and restart the SSR management process on the M node. Repeat this process on each M node, one at a time.
2. Stop the SSR management process and RADIUS process on the SM node, upgrade the SM node (install and configure the new software), and restart the SSR management process and RADIUS process on the SM node. Repeat this process on each SM node, one at a time.
3. Stop the RADIUS process on the S node, install the new software on the S node, configure the new software, and restart the SSR data process on the S node.
4. Stop the SSR data process on the D node, install the new software on the D node, configure the new software, and restart the SSR data process on the D node. Repeat this process on each D node, one at a time.

## Introduction and Requirements

Using these instructions, you will be able to upgrade your working SSR cluster from any of the specified software releases without interruption of service. However, because the cluster will momentarily run with one node lesser than normally required, it is advisable to schedule the maintenance window outside any busy hours.

**NOTE:** Skipping versions when upgrading the cluster using the rolling restart method is not supported. Since SBR Carrier 8.0.0 uses MySQL 5.5.37, and 8.4.0, 8.4.1, and 8.5.0 use 5.7.18 for Linux (see [Table 47 on page 398](#)), we strongly recommend that you not use the rolling restart method to upgrade the cluster version of SBR Carrier directly from release 8.0.0 to 8.4.x or later on a Linux machine. On Solaris, we strongly recommend that you not use the rolling restart method to upgrade the cluster version of SBR Carrier directly from release 8.0.0 to 8.6.0 or later, since SBR Carrier 8.0.0 uses MySQL 5.5.37 and 8.6.0 uses 5.7.25 for Solaris. Instead, use the backup, destroy, and re-create method to upgrade or perform a clean install.

**Table 47: MySQL and NDB Versions Used by SBR Carrier**

SBR Carrier Version	MySQL Version	NDB Version
8.0.0	5.5.37	7.2.16
8.1.0	5.6.22	7.3.8
8.2.0	5.6.28	7.4.10
8.3.0	5.6.29	7.4.11
8.4.0, 8.4.1, and 8.5.0	Linux: 5.7.18 Solaris: 5.6.36	Linux: 7.5.6 Solaris: 7.4.15
8.6.0	5.7.25	7.6.9
Starting from 8.6.0R14 version for Solaris 11.4.25.0.1.75.3 and RHEL 8.1	8.0.22	8.0.22

**NOTE:** Upgrades from SBR Carrier releases earlier than 7.5.0 have not been tested by Juniper Networks.

**NOTE:** The entire upgrade process takes approximately three hours to complete. The current number of concurrent sessions in SSR and the current load on the data nodes contribute to the time it takes to complete this upgrade.

To perform the rolling restart upgrade, you will need the following:



- The original cluster configuration files from the previous release of SBR Carrier installation (the contents of `/opt/JNPRshare`).
- The SBR Carrier cluster distribution files, for example: **sbr-cl-8.6.0.R-1.sparc.tgz**.
- At least 10 GB of disk space on each machine running an SBRC front-end (S or SM node).
- At least 3 GB of disk space on each machine running a data node (D node).

## Preparation

Complete the following steps to prepare for the upgrade:

1. Verify that all your front-end S or SM nodes are configured with matching keys in the **spi.ini** file. For details, see the **spi.ini** file section in the *SBR Carrier Reference Guide*, and ensure that all your S or SM nodes have the same keys and are listed in the **[Hosts]** section.



**CAUTION:** Do not proceed with this upgrade until all your S or SM nodes are configured with matching keys and are listed in the **[Hosts]** section of the **spi.ini** file. Failure to do so will cause the upgrade to fail.

2. Copy the **sbr-cl-8.6.0.R-1.sparc.tgz** software distribution package to every node in the cluster, regardless of its role (S, SM, D, and M nodes). Copy the package to the `/opt/tmp` directory or another location with enough disk space, and unzip and un-tar the distribution package.

For example:

```
cd /opt/tmp
```

```
gunzip sbr-cl-8.6.0.R-1.sparc.tgz
```

```
tar xf sbr-cl-8.6.0.R-1.sparc.tar
```

3. (Optional) For simplicity, in this procedure the entire distribution is copied to every node in the cluster. If you have limited disk space, you can optionally just copy the file **\*mysqlcluster\*.tar.gz** to every data (D) node and management (M) node, or mount it over NFS. This is the only file that is actually required on data nodes and management (M) nodes. If you are comfortable with Solaris, you can choose to unpack the entire distribution only on the S and SM nodes, and copy just **\*mysqlcluster\*.tar.gz** to the M and D nodes. If you choose this advanced method, you must be sure to unpack the entire distribution on all S or SM nodes. Note the path to the file **\*mysqlcluster\*.tar.gz** on your D nodes, for example: `/opt/tmp/JNPRsbr.pkg/reloc/radius/install/ndb/bin/*mysqlcluster*.tar.gz`.

**NOTE:** When SBR Carrier is upgraded from a release earlier than 7.4.0 (inclusive) to 7.4.1 or a later release, the `/opt/tmp/JNPRsbr/pkg/reloc/radius/install/ndb/bin/` path contains two versions of MySQL files. We recommend that you use the latest version of MySQL file.

4. Verify that you have more than 6 GB of disk space remaining on each node running SBR (S or SM nodes) after copying the entire distribution package.
5. Verify that you have more than 1 GB of disk space remaining on each D node after copying the package.
6. Verify that you have more than 1 GB of disk space remaining on each M node after copying the package.

## Upgrading the M Nodes in the Cluster

To perform a rolling restart upgrade in your cluster, first you need to install the new software and configure it on each M node in the cluster.

**NOTE:** Execute this procedure on each M node, one at a time. This example procedure uses two M nodes `bng-mars.englab.juniper.net` (10.212.10.68) and `bng-sbr-perf1` (10.212.10.66).

**NOTE:** If you have limited disk space, you can just copy the file `*mysqlcluster*.tar.gz` to every M node, or mount it over NFS. Refer to [“Installing Only the Latest MySQL Package” on page 435](#).

### Installing the New SBRC Software on the M Nodes

Execute the following steps to install the new SBRC software package on the M node:

1. Stop the SSR process.

Log in to the node as root and execute:

**root@bng-mars: /opt/JNPRsbr/radius/sbrd stop ssr**

```
Stopping ssr auxiliary processes
Stopping ssr management processes
Connected to Management Server at: 10.212.10.68:5235
```

```
Node 52 has shutdown.
Disconnecting to allow Management Server to shutdown
```

2. Log in to the SM node and execute **./sbrd status**. Verify that the M node is disconnected from other nodes.

Execute:

**root@sbr1: /opt/JNPRsbr/radius/sbrd status**

```
-----
SBR 7.60-R1.50242 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_120011-14 node sbr1.englab.juniper.net(sm)
-----

Connected to Management Server at: 10.212.10.213:5235

[ndbd(NDB)] 2 node(s)
id=1 @10.212.10.188 (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0, Master)
id=2 @10.212.10.67 (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0)

[ndb_mgmd(MGM)] 3 node(s)
id=51 @10.212.10.213 (mysql-5.5.20 ndb-7.2.5)
id=52 (not connected, accepting connect from 10.212.10.68)
id=53 @10.212.10.66 (mysql-5.5.20 ndb-7.2.5)

[mysqld(API)] 8 node(s)
id=61 @10.212.10.213 (mysql-5.5.20 ndb-7.2.5)
id=62 (not connected, accepting connect from 10.212.10.68)
id=63 @10.212.10.66 (mysql-5.5.20 ndb-7.2.5)
id=100 @10.212.10.213 (mysql-5.5.20 ndb-7.2.5)
id=101 @10.212.10.65 (mysql-5.5.20 ndb-7.2.5)

hadm      25074 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
hadm      25141 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
hadm      25462 /opt/JNPRmysql/install/bin/mysqld
--basedir=/opt/JNPRmysql/install --datadir=/opt/JNPRmysqld/data
--plugin-dir=/usr/local/mysql/lib/plugin
--log-error=/opt/JNPRmysqld/mysqld_safe.err
--pid-file=/opt/JNPRmysqld/mysqld.pid --socket=/opt/JNPRhadm/.mysql.sock
--port=3001
root      27191 radius sbr.xml
```

3. Navigate to the **install** directory.

Execute:

**root@bng-mars: /opt/JNPRsbr/radius/install**

#### 4. Unconfigure the node.

Execute:

**root@bng-mars:./unconfigure**

```
Unconfiguring SBR Software
```

```
-----
SBR 7.60-R1.50242 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_120011-14 node bng-mars.englab.juniper.net(m)
is configured and processes are down, may be reconfigured if desired
-----
```

```
Unconfiguring Cluster Node...
```

```
WARNING: You are about to unconfigure this node.
Are you sure that you wish to continue? (y,n): y
```

```
Cleaning directories
```

```
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRmysqld
/opt/JNPRndb_mgmd
/opt/JNPRndbd
```

```
Locating shared directory...
```

```
drwxrwxr-x   2 hadm      hadmg      512 Jan 11 13:07
/opt/JNPRshare/install/6node_cluster
```

```
WARNING: If you remove the shared directory for this cluster, you
will either have to recover the data from another cluster node or
reconfigure the entire cluster again. This is neither necessary
nor recommended if you are updating an existing configuration.
Remove the shared directory for this cluster? [n]:
```

```
Locating OS user account and home directory...
```

```
hadm:x:50029:65541:SBR software:/opt/JNPRhadm:/bin/bash
hadmg::65541:
drwxrwx---   2 hadm      hadmg      1536 Jan 11 14:30 /opt/JNPRhadm
```

```
WARNING: If you remove the OS user account hadm you will have
to recreate it, the associated OS group account hadmg, and
```

```
the associated home directory /opt/JNPRhadm
This is neither necessary nor recommended if you are updating
an existing configuration.  Remove the OS user account? [n]:
```

```
Unconfigured
```

5. Move your original installation to a different location to prevent the removal of files. Be sure not to move the **/opt/JNPRshare** directory, which should contain your original cluster configuration files. For example, if you previously installed Release 7.6.0, execute the following as root:
  - a. Navigate to the opt directory: **root@bng-mars: cd /opt.**
  - b. Create a new directory: **root@bng-mars: mkdir 760.**
  - c. Move the configuration files to the new directory: **root@bng-mars: mv JNPRsbr JNPRmysql JNPRmysqld JNPRhadm JNPRndb\_mgmd JNPRndbd 760.**
  - d. Remove the old software package: **root@bng-mars: pkgrm JNPRsbr.**
6. Install the new SBRC software distribution package.

Execute:

```
root@bng-mars: cd /opt/tmp pkgadd -d
```

```
The following packages are available:
  1  JNPRsbr.pkg      JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier
Cluster Edition)

                        (sparc) 8.60.50242

Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,??,q]: 1
```

7. Enter **1** to install the **JNPRsbr.pkg**. The system installs the new software.

```
Processing package instance JNPRsbr.pkg
.
.
.
Installation of JNPRsbr.pkg was successful.
```

8. Once the system displays the message **Installation of JNPRsbr.pkg was successful**, repeat this procedure on the next SM node. Once the new software is installed on each SM node, proceed to [“Configuring the New SBRC Software on the M Node” on page 404.](#)

## Configuring the New SBRC Software on the M Node

Configure the new software on each M node, one at a time:

1. Configure the SBRC software on the M node, using the original cluster configuration files.

Execute:

**root@bng-mars: ./configure**

```
Configuring SBR Software
```

```
END USER LICENSE AGREEMENT READ THIS END USER LICENSE AGREEMENT ( "AGREEMENT" ).....
```

```
.  
.  
.
```

```
Do you accept the terms in the license agreement? [n]:    y
```

2. Enter **y** to accept the license agreement.

```
-----  
SBR 8.60-R1.50765 cluster  
on SunOS 5.10 Generic_141444-09 node bng-mars.englab.juniper.net  
is not configured and processes are down, needs to be configured  
-----
```

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.
2. Generate Cluster Definition  
Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.
3. Configure Cluster Node  
To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.
4. Reconfigure RADIUS Server  
Only on SBR nodes, updates the existing SBR configuration.
5. Create Temporary Cluster

```
Used to approximate a cluster using only this one machine.
Intended for migration and demonstration purposes only.
```

#### 6. Upgrade From Restricted Cluster License

```
Used to upgrade from restricted cluster to regular cluster.
Removes database restriction on the number of concurrent
sessions and enables the addition of an expansion kit license
```

```
Enter the number of the desired configuration task or quit (2,q): 3
```

### 3. Enter **3** to select option **3. Configure Cluster Node**.

```
-----
SBR 8.60-R1.50765 cluster
on SunOS 5.10 Generic_141444-09 node bng-mars.englab.juniper.net
is not configured and processes are down, needs to be configured
-----

Configuring Cluster Node...
```

### 4. Specify the name of your existing cluster. In this example the cluster name is "6node\_cluster".

```
Enter SBR cluster name [bng-mars]: 6node_cluster
```

```
Reading shared configuration from /opt/JNPRshare/install/6node_cluster
.....
```

### 5. Migrate the configuration from your previous software release to the new software package. In most cases, you simply need to identify your previous installation. In this example, the cluster configuration is named **configure.6node\_cluster.tar** (see the previous step), and is stored in the **/opt/JNPRsbr/radius/install** directory. Once you specify the configuration, SBRC migrates your configuration.

Enter:

```
root@bng-mars: /opt/JNPRsbr/radius/install
```

```
Reading shared configuration from /opt/JNPRshare/install/6node_cluster

Enter absolute path to configure.6node_cluster.tar
```

```
if you have it, or quit (q) so that you can get it from
/opt/JNPRsbr/radius/install on the first cluster node.
```

```
Generating configuration files
```

```
Reviewing configuration files
```

```
/opt/JNPRsbr/radius/install/tmp/config.ini
```

```
/opt/JNPRsbr/radius/install/tmp/my.cnf
```

```
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
```

```
View (v), accept (a), or reject (r) configuration files: a
```

6. Enter **a** to accept the configuration files.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y
```

7. Enter **y** to continue.

```
Cleaning directories
```

```
/opt/JNPRhadm
```

```
/opt/JNPRmysql
```

```
/opt/JNPRmysqld
```

```
/opt/JNPRndb_mgmd
```

```
/opt/JNPRndbd
```

```
Applying configuration
```

```
Initializing Session State Register, please wait a few minutes...
```

```
-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node bng-mars.englab.juniper.net(m)
is configured and processes are down, may be reconfigured if desired
-----
```

#### 1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

#### 2. Generate Cluster Definition

Creates new or updates existing cluster definitions.



Modifies the shared directory but does not modify this node.

3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.

4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.

6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

READY: last operation succeeded, node configured.

Enter the number of the desired configuration task or quit (4,q): q

8. Start the SSR process on the M node.

Execute:

**root@bng-mars: /opt/JNPRsbr/radius/sbrd start ssr**

```
Starting ssr management processes
MySQL Cluster Management Server mysql-5.7.25 ndb-7.6.9
2016-01-11 15:47:33 [MgmtSrvr] INFO      -- Skipping check of config directory
since config cache is disabled.
-Node 52: connected (Version 7.4.10)

Starting ssr auxiliary processes
```

9. Check the status of the M node and verify it has properly reconnected to the cluster.

Execute:

**root@bng-mars: /opt/JNPRsbr/radius/sbrd status**

```
-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node bng-mars.englab.juniper.net(m)
```

```

-----

Connected to Management Server at: 10.212.10.68:5235

[ndbd(NDB)]      2 node(s)
id=1    @10.212.10.188  (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0, *)
id=2    @10.212.10.67  (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0)

[ndb_mgmd(MGM)]  3 node(s)
id=51   @10.212.10.213  (mysql-5.5.20 ndb-7.2.5)
id=52   @10.212.10.68  (mysql-5.6.36 ndb-7.4.15)
id=53   @10.212.10.66  (mysql-5.5.20 ndb-7.2.5)

[mysqld(API)]    5 node(s)
id=61   @10.212.10.213  (mysql-5.5.20 ndb-7.2.5)
id=62   @10.212.10.68  (mysql-5.6.36 ndb-7.4.15)
id=63   @10.212.10.66  (mysql-5.5.20 ndb-7.2.5)
id=100  @10.212.10.213  (mysql-5.5.20 ndb-7.2.5)
id=101  @10.212.10.65  (mysql-5.5.20 ndb-7.2.5)

-----

Current state of network interfaces:

-----

    hadm 15410 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
    hadm 15426 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
    hadm 15615 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
    --datadir=/o

```

- 10.(Optional) Make any modifications to the **config.ini** file in the **/opt/JNPRhadm** directory. For example, you may need to make adjustments to **DataMemory** based on the previous value recommended for your existing cluster.

Once you make any adjustments to the **config.ini** file, you must restart each node hosting a management server (SM and M nodes).

Execute:

root@bng-mars:/opt/JNPRsbr/radius/sbrd restart

**NOTE:** All SM and M nodes must have the same values in the **config.ini** file.

11. Verify that SBRC is performing correctly by checking logs and counters.
12. Repeat this process on the other M node; see [“Upgrading the M Nodes in the Cluster” on page 400](#).

## Upgrading the SM Nodes in the Cluster

After all M nodes in the cluster are upgraded, you need to upgrade the software in SM nodes. To perform a rolling restart upgrade in SM nodes of your cluster, install the new software and configure it on each SM node in the cluster.

**NOTE:** Execute this procedure on each SM node, one at a time. This example procedure uses one SM node **sbr1.englab.juniper.net** (10.212.10.213).

**NOTE:** If you have limited disk space, you can just copy the file **\*mysqlcluster\*.tar.gz** to every SM node, or mount it over NFS. Refer to [“Installing Only the Latest MySQL Package” on page 435](#).

### Installing the New SBRC Software on the SM Node

Execute the following steps to install the new SBRC software package on the SM node:

1. Stop the SSR process.

Log in to the node as root and execute:

```
root@sbr1: /opt/JNPRsbr/radius/sbrd stop ssr
```

```
Stopping ssr auxiliary processes
Stopping ssr management processes
Connected to Management Server at: 10.212.10.213:5235
Node 51 has shutdown.
Disconnecting to allow Management Server to shutdown
```

2. Log in to a M node and execute **./sbrd status**. Verify that the SM node is disconnected from the other nodes.

Execute:

```
root@bng-mars: /opt/JNPRsbr/radius/sbrd status
```

```

-----
SBR 7.60-R1.50242 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_120011-14 node bng-mars.englab.juniper.net(m)
-----

Connected to Management Server at: 10.212.10.66:5235

[ndbd(NDB)] 2 node(s)
id=1 @10.212.10.188 (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0, *)
id=2 @10.212.10.67 (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0)

[ndb_mgmd(MGM)] 3 node(s)
id=51 (not connected, accepting connect from 10.212.10.213)
id=52 @10.212.10.68 (mysql-5.6.36 ndb-7.4.15)
id=53 @10.212.10.66 (mysql-5.6.36 ndb-7.4.15)

[mysqld(API)] 5 node(s)
id=61 (not connected, accepting connect from 10.212.10.213)
id=62 @10.212.10.68 (mysql-5.6.36 ndb-7.4.15)
id=63 @10.212.10.66 (mysql-5.6.36 ndb-7.4.15)
id=100 (not connected, accepting connect from 10.212.10.213)
id=101 @10.212.10.65 (mysql-5.5.20 ndb-7.2.5)

hadm      17775 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
hadm      17849 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
hadm      18196 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install

--datadir=/opt/JNPRmysqld/data --plugin-dir=/opt/JNPRmysql/install/lib/plugin
--log-error=/opt/JNPRmysqld/mysqld_safe.err
--pid-file=/opt/JNPRmysqld/mysqld.pid --socket=/opt/JNPRhadm/.mysql.sock
--port=3001

```

### 3. Stop the RADIUS process in the SM node.

Execute:

**root@sbr1: /opt/JNPRsbr/radius/sbrd stop radius**

```

Stopping radius server processes
waiting for radius
10 seconds elapsed, still waiting
radius stopped

```

Check the status of the SM node after stopping the RADIUS process.

Execute:

**root@sbr1: /opt/JNPRsbr/radius/sbrd status**

```
-----
SBR 7.60-R1.50242 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_120011-14 node sbr1.englab.juniper.net(sm)
-----
```

4. Navigate to the **install** directory.

Execute:

**root@sbr1: /opt/JNPRsbr/radius/install**

5. Unconfigure the node.

Execute:

**root@sbr1:./unconfigure**

Unconfiguring SBR Software

```
-----
SBR 7.60-R1.50242 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_120011-14 node sbr1.englab.juniper.net(sm)
is configured and processes are down, may be reconfigured if desired
-----
```

Unconfiguring Cluster Node...

WARNING: You are about to unconfigure this node.  
Are you sure that you wish to continue? (y,n): y

Cleaning directories

```
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRmysqld
/opt/JNPRndb_mgmd
/opt/JNPRndbd
```

Locating shared directory...

```
drwxrwxr-x  2 hadm      hadmg      512 Jan 11 02:32
/opt/JNPRshare/install/6node_cluster
```

WARNING: If you remove the shared directory for this cluster, you will either have to recover the data from another cluster node or reconfigure the entire cluster again. This is neither necessary

```

nor recommended if you are updating an existing configuration.
Remove the shared directory for this cluster? [n]:

Locating OS user account and home directory...
hadm:x:50025:65540:SBR software:/opt/JNPRhadm:/bin/bash
hadmg::65540:
drwxrwx---  2 hadm      hadmg      1536 Jan 11 06:54 /opt/JNPRhadm

WARNING: If you remove the OS user account hadm you will have
to recreate it, the associated OS group account hadmg, and
the associated home directory /opt/JNPRhadm
This is neither necessary nor recommended if you are updating
an existing configuration. Remove the OS user account? [n]:

Unconfigured

```

6. Move your original installation to a different location to prevent the removal of files. Be sure not to move the **/opt/JNPRshare** directory, which should contain your original cluster configuration files. For example, if you previously installed Release 7.6.0, execute the following as root:
  - a. Navigate to the opt directory: **root@sbr1: cd /opt.**
  - b. Create a new directory: **root@sbr1: mkdir 760.**
  - c. Move the configuration files to the new directory: **root@sbr1: mv JNPRsbr JNPRmysql JNPRmysqld JNPRhadm JNPRndb\_mgmd JNPRndbd 760.**
  - d. Remove the old software package: **root@sbr1: pkgrm JNPRsbr.**
7. Install the new SBRC software distribution package.

Execute:

```
root@sbr1: cd /opt/tmp pkgadd -d
```

```

The following packages are available:
  1  JNPRsbr.pkg      JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier
Cluster Edition)
                                   (sparc) 8.60.50242

Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,??,q]: 1

```

8. Enter **1** to install the **JNPRsbr.pkg**. The system installs the new software.

```
Processing package instance JNPRsbr.pkg
.
.
.
Installation of JNPRsbr.pkg was successful.
```

9. Once the system displays the message **Installation of JNPRsbr.pkg was successful**, repeat this procedure on the next SM node. Once the new software is installed on each SM node, proceed to [“Configuring the New SBRC Software on the SM Node” on page 413](#).

## Configuring the New SBRC Software on the SM Node

Configure the new software on each SM node, one at a time:

1. Configure the SBRC software on the SM node, using the original cluster configuration files.

Execute:

**root@sbr1: ./configure**

```
Configuring SBR Software

END USER LICENSE AGREEMENT READ THIS END USER LICENSE AGREEMENT ("AGREEMENT")
.....
.
.
.
.Do you accept the terms in the license agreement? [n]:    y
```

2. Enter **y** to accept the license agreement.

```
-----
SBR 8.60-R1.50765 cluster
on SunOS 5.10 Generic_141444-09 node sbr1.englab.juniper.net
is not configured and processes are down, needs to be configured
-----

1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.

2.  Generate Cluster Definition
```

```

Creates new or updates existing cluster definitions.
Modifies the shared directory but does not modify this node.

3.  Configure Cluster Node
    To be preceded by 'Generate Cluster Definition' on one node.
    Must be invoked on each and every node of the cluster.

4.  Reconfigure RADIUS Server
    Only on SBR nodes, updates the existing SBR configuration.

5.  Create Temporary Cluster
    Used to approximate a cluster using only this one machine.
    Intended for migration and demonstration purposes only.

6.  Upgrade From Restricted Cluster License
    Used to upgrade from restricted cluster to regular cluster.
    Removes database restriction on the number of concurrent
    sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q): 3

```

3. Enter **3** to select option **3. Configure Cluster Node**.

```

-----
SBR 8.60-R1.50765 cluster
on SunOS 5.10 Generic_141444-09 node sbr1.englab.juniper.net
is not configured and processes are down, needs to be configured
-----

Configuring Cluster Node...

```

4. Specify the name of your existing cluster. In this example the cluster name is “6node\_cluster”.

```

Enter SBR cluster name [sbr1]: 6node_cluster

```

```

Reading shared configuration from /opt/JNPRshare/install/6node_cluster
.....

```

5. Migrate the configuration from your previous software release to the new software package. In most cases, you simply need to identify your previous installation. In this example, the cluster configuration



is named **configure.6node\_cluster.tar** (see the previous step), and is stored in the **/opt/JNPRsbr/radius/install** directory. Once you specify the configuration, SBRC migrates your configuration.

Enter:

**root@sbr1: /opt/JNPRsbr/radius/install**

```
Reading shared configuration from /opt/JNPRshare/install/6node_cluster
```

```
Enter absolute path to configure.6node_cluster.tar
if you have it, or quit (q) so that you can get it from
/opt/JNPRsbr/radius/install on the first cluster node.
```

```
Generating configuration files
```

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files: a
```

6. Enter **a** to accept the configuration files.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y
```

7. Enter **y** to continue.

```
Cleaning directories
```

```
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRmysqld
/opt/JNPRndb_mgmd
/opt/JNPRndbd
```

```
Applying configuration
```

```
Initializing Session State Register, please wait a few minutes...
19949 Killed
```

```

Do you want to configure Java Runtime Environment for JDBC Feature [n] :

Please enter backup or radius directory from which to migrate.
Enter n for new configuration, s to search, or q to quit
[/opt/JNPRsbr/radius/install/backups/2010:10:21-08:31:21]:

Migrating files, please wait...
-

WARNING: A manual merge may be required between the files located in
/opt/JNPRsbr/radius (working files) and
/opt/JNPRsbr/radius/install/backups/2010:10:21-08:32:18 (shipped files)

```

8. Press **Enter** through the following prompts to keep your configuration.

```

Enter initial admin user (UNIX account must have a valid password) [root]:
Configuring for use with generic database
Do you want to enable "Radius WatchDog" Process? [n]:
Do you want to enable LCI? [n]:
Do you want to configure for use with Oracle? [n]:
Removing oracle references from startup script
Do you want to configure for use with SIGTRAN? [n]:
Removing SIGTRAN references from startup script
Do you want to configure SNMP? [n]:
Configuring Admin GUI Webserver
Compatible Java version found :
Do you want to install custom SSL certificate for Admin WebServer? [n]:

```

9. Enter **e** when prompted to enable the autoboot scripts.

```

Enable (e), disable (d), or preserve (p) autoboot scripts [e]: e

```

10. The main configuration menu is displayed.

Enter **q** to quit.

```

The SBR Admin Web GUI can be launched using the following URL:
https://<servername>:2909

-----

SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node sbr1.englab.juniper.net(sm)

```

is configured and processes are down, may be reconfigured if desired

```
-----

1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.

2.  Generate Cluster Definition
    Creates new or updates existing cluster definitions.
    Modifies the shared directory but does not modify this node.

3.  Configure Cluster Node
    To be preceded by 'Generate Cluster Definition' on one node.
    Must be invoked on each and every node of the cluster.

4.  Reconfigure RADIUS Server
    Only on SBR nodes, updates the existing SBR configuration.

5.  Create Temporary Cluster
    Used to approximate a cluster using only this one machine.
    Intended for migration and demonstration purposes only.

6.  Upgrade From Restricted Cluster License
    Used to upgrade from restricted cluster to regular cluster.
    Removes database restriction on the number of concurrent
    sessions and enables the addition of an expansion kit license

READY: last operation succeeded, node configured.
Enter the number of the desired configuration task or quit (4,q): q
```

11. Start the SSR process on the SM node.

Execute:

**root@sbr1: /opt/JNPRsbr/radius/sbrd start ssr**

```
Starting ssr management processes
MySQL Cluster Management Server mysql-5.7.25 ndb-7.6.9
2016-01-11 07:18:49 [MgmtSrvr] INFO      -- Skipping check of config directory
since config cache is disabled.
-Node 51: connected (Version 7.4.10)

Starting ssr auxiliary processes
```

12. Log in as the user 'hadm', navigate to the hadm user's home directory (/opt/JNPRhadm/ by default), and re-create the stored procedures by executing **CreateDB.sh**.

If you need to customize the sessions database to match your existing cluster session database, see [“Customizing the SSR Database Current Sessions Table” on page 144](#). Any customization must be done prior to running the **CreateDB.sh** script.

**NOTE:** Despite its name, this command does not change the contents of your database, since the database already exists. However, it re-creates the stored procedures that are required by the administration scripts. Failure to execute this step results in inoperable management scripts, such as **ShowSessions.sh**, **DelSession.sh**, and **CreateDB.sh**.

### 13. Execute the **UpdateSchema.pl** utility for schema upgrades.

- a. First, you need to determine whether your schema needs to be updated for this release by executing **UpdateSchema** with no arguments as the hadm user. This command returns a list of schemas which need to be updated.

Execute:

```
hadm@sbr1:~>perl ./UpdateSchema.pl
```

```
Usage: perl UpdateSchema.pl [version]
      where [version] is 5.4, 5.5, 7.2.5, or 8.4 and above
      Other versions do not need a schema update.
```

From SBR 8.5 release, to increase the Column size of AcctSessionId from 24 to 48, the below usage is recommended:

```
Usage: perl UpdateSchema.pl <To_Version> ColumnUpdate:Sbr_AcctSessionId
```

**NOTE:** Future SBR Carrier software releases may require additional updates, we recommend you execute this step every time you upgrade to determine if a schema upgrade is required.

- b. If you are running a software release that requires a schema change, execute the command again with the version to which you are upgrading. For example, if you are upgrading to SBR Carrier Release 8.6.0:

Execute:

```
hadm@sbr1:~>perl ./UpdateSchema.pl 8.6 ColumnUpdate:Sbr_AcctSessionId
```

```
Altering Tables for adding Sbr_Ipv6Address- Please wait...Done
Altering Tables for changing Sbr_AcctSessionId size from 24 to 48- Please
wait... Done
```

**NOTE:** The time required for altering the CST depends on the number of sessions stored in the CST. For example, SBR Carrier takes approximately 30 minutes to update a CST that contains 5 million sessions.

- c. In SBR 8.6.0 R15 RHEL 8.1 and Solaris 11.4.25.0.1.75.3 variants, the size of "Name" column which represents the Pool Name present in the table "Sbr\_lpPools" is modified from 24 varchar to 84 varchar.

To accommodate and synchronize this change between the existing and planned upgrade SBR version (SBR 8.6.0 R15 on RHEL 8.1 and Solaris 11.4.25.0.1.75.3), the below command **MUST** be executed only once on the first upgrade planned cluster node.

The following command should be executed from "hadm" mode after successful installation and configuration of the Node.

```
$ su - hadmhadm@bng-perf5-kvm5:~> perl ./UpdateSchema.pl 8.6 ColumnUpdate:Name
Schema is already updated with Sbr_ipv6Address column
Altering Tables(Sbr_lpPools) to increase Name size from 24 to 84
Please wait... Done
```

#### 14. Start the RADIUS process on the SM node.

Execute:

**root@sbr1:/opt/JNPRsbr/radius/sbrd start radius**

```
Starting radius server processes
RADIUS: Process ID of daemon is 5405
RADIUS: Starting DCF system
RADIUS: Configuration checksum: 6A 3B 7A D5

root@sbr1 [/opt/JNPRsbr/radius]
$Starting webserver
Document Root: /opt/JNPRsbr/radius/
. 2016-01-11 07:21:17.217:INFO::main: Logging initialized @4945ms
. . . . . Webserver started

root@sbr1 [/opt/JNPRsbr/radius]
$AppleTalk
AppleTalk
AppleTalk
RADIUS:
RADIUS: Steel-Belted Radius licenses
```

```

RADIUS: License String:                               Additional Info:
RADIUS: XXXXXXXXXXXXXXXXXXXX                         SSR Cluster & Management Expansion
RADIUS: XXXXXXXXXXXXXXXXXXXX                         SSR Management Expansion
RADIUS: XXXXXXXXXXXXXXXXXXXX                         Run license - 50k sessions
RADIUS:
RADIUS: Total concurrent sessions licensed: 50000
RADIUS:
RADIUS: Unlimited Transactions Per Second is Allowed
RADIUS:
RADIUS: Licenses enable this server to:
RADIUS:         Run
RADIUS:         use Multi Node Cluster
RADIUS:         connect '3' management and '2' cluster nodes to database
RADIUS:
RADIUS: Licensed for Steel-Belted Radius Carrier

root@sbr1 [/opt/JNPRsbr/radius]
$RADIUS: Attribute Editing enabled
RADIUS: DCF system started

```

15. Check the status of the SM node and verify it has properly reconnected to the cluster.

Execute:

**root@sbr1: /opt/JNPRsbr/radius/sbrd status**

```

-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node sbr1.englab.juniper.net(sm)
-----

Connected to Management Server at: 10.212.10.213:5235

[ndbd(NDB)]      2 node(s)
id=1   @10.212.10.188  (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0, *)
id=2   @10.212.10.67  (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0)

[ndb_mgmd(MGM)]  3 node(s)
id=51   @10.212.10.213  (mysql-5.6.36 ndb-7.4.15)
id=52   @10.212.10.68  (mysql-5.6.36 ndb-7.4.15)
id=53   @10.212.10.66  (mysql-5.6.36 ndb-7.4.15)

[mysqld(API)]    5 node(s)
id=61   @10.212.10.213  (mysql-5.6.36 ndb-7.4.15)

```

```

id=62   @10.212.10.68   (mysql-5.6.36 ndb-7.4.15)
id=63   @10.212.10.66   (mysql-5.6.36 ndb-7.4.15)
id=100  @10.212.10.213  (mysql-5.6.36 ndb-7.4.15)
id=101  @10.212.10.65   (mysql-5.5.20 ndb-7.2.5)

-----

Current state of network interfaces:

10.212.10.213.1646          Idle
10.212.10.213.1813          Idle
10.212.10.213.1645          Idle
10.212.10.213.1812          Idle
      *.1812                *.*          0      0 49152      0 LISTEN
      *.1813                *.*          0      0 49152      0 LISTEN

-----

hadm  4761 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
hadm  4779 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
hadm  4968 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
--datadir=/o
root  5405 radius sbr.xml
root  5475 webserver

```

- 16.(Optional) Make any modifications to the **config.ini** file in the **/opt/JNPRhadm** directory. For example, you may need to make adjustments to **DataMemory** based on the previous value recommended for your existing cluster.

Once you make any adjustments to the **config.ini** file, you must restart each node hosting a management server (SM nodes).

Execute:

root@sbr1:/opt/JNPRsbr/radius/sbrd restart

**NOTE:** All SM nodes must have the same values in the **config.ini** file.

17. Verify that SBRC is performing correctly by checking logs and counters.

18. Repeat this process on the other SM node; see [“Upgrading the SM Nodes in the Cluster” on page 409](#).

## Upgrading the S Nodes in the Cluster

To perform a rolling restart upgrade in S nodes of your cluster, install the new software and configure it on each S node in the cluster.

**NOTE:** Execute this procedure on each S node, one at a time. This example procedure uses one S node bng-sbrha-1 (10.212.10.65).

### Installing the New SBRC Software on the S Node

Execute the following steps to install the new SBRC software package on the S node:

1. Stop the RADIUS process in the S node.

Execute:

**root@bng-sbrha-1: /opt/JNPRsbr/radius/sbrd stop radius**

```
Stopping radius server processes
waiting for radius
10 seconds elapsed, still waiting
radius stopped
```

Log in to a M node and execute **./sbrd status**. Verify that the S node is disconnected from the other nodes.

Execute:

**root@bng-mars: /opt/JNPRsbr/radius/sbrd status**

```
-----
SBR 7.60-R1.50242 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_120011-14 node bng-mars.englab.juniper.net(m)
-----
```

```
Connected to Management Server at: 10.212.10.68:5235
```

```
[ndbd(NDB)] 2 node(s)
id=1 @10.212.10.188 (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0, *)
id=2 @10.212.10.67 (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 3 node(s)
```



```

id=51 @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=52 @10.212.10.68 (mysql-5.6.36 ndb-7.4.15)
id=53 @10.212.10.66 (mysql-5.6.36 ndb-7.4.15)

[mysqld(API)] 5 node(s)
id=61 @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=62 @10.212.10.68 (mysql-5.6.36 ndb-7.4.15)
id=63 @10.212.10.66 (mysql-5.6.36 ndb-7.4.15)
id=100 @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=101 (not connected, accepting connect from 10.212.10.65)

hadm      4354 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
hadm      4908 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
hadm      5255 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install

--datadir=/opt/JNPRmysqld/data --plugin-dir=/opt/JNPRmysql/install/lib/plugin
--log-error=/opt/JNPRmysqld/mysqld_safe.err
--pid-file=/opt/JNPRmysqld/mysqld.pid --socket=/opt/JNPRhadm/.mysql.sock
--port=3001

```

## 2. Navigate to the **install** directory.

Execute:

```
root@bng-sbrha-1: /opt/JNPRsbr/radius/install
```

## 3. Unconfigure the node.

Execute:

```
root@bng-sbrha-1:./unconfigure
```

```

Unconfiguring SBR Software

-----

SBR 7.60-R1.50242 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_120011-14 node bng-sbrha-1(s)
is configured and processes are down, may be reconfigured if desired
-----

Unconfiguring Cluster Node...

WARNING: You are about to unconfigure this node.
Are you sure that you wish to continue? (y,n): y

Cleaning directories
/opt/JNPRhadm

```

```

Locating shared directory...
drwxrwxr-x   2 hadm      hadmg          512 Jan 11 13:25
/opt/JNPRshare/install/6node_cluster

WARNING: If you remove the shared directory for this cluster, you
will either have to recover the data from another cluster node or
reconfigure the entire cluster again. This is neither necessary
nor recommended if you are updating an existing configuration.
Remove the shared directory for this cluster? [n]:

Locating OS user account and home directory...
hadm:x:554542:554539:SBR software:/opt/JNPRhadm:/bin/bash
hadmg::554539:
drwxrwx---   2 hadm      hadmg          512 Jan 11 19:47 /opt/JNPRhadm

WARNING: If you remove the OS user account hadm you will have
to recreate it, the associated OS group account hadmg, and
the associated home directory /opt/JNPRhadm
This is neither necessary nor recommended if you are updating
an existing configuration. Remove the OS user account? [n]:

Unconfigured

```

4. Move your original installation to a different location to prevent the removal of files. Be sure not to move the **/opt/JNPRshare** directory, which should contain your original cluster configuration files. For example, if you previously installed Release 7.6.0, execute the following as root:
  - a. Navigate to the opt directory: **root@bng-sbrha-1: cd /opt.**
  - b. Create a new directory: **root@bng-sbrha-1: mkdir 760.**
  - c. Move the configuration files to the new directory: **root@bng-sbrha-1: mv JNPRsbr JNPRmysql JNPRmysqld JNPRhadm JNPRndb\_mgmd JNPRndbd 760.**
  - d. Remove the old software package: **root@bng-sbrha-1: pkgrm JNPRsbr.**
5. Install the new SBRC software distribution package.

Execute:

**root@bng-sbrha-1: cd /opt/tmp pkgadd -d**

```

The following packages are available:
  1  JNPRsbr.pkg      JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier
Cluster Edition)
                                (sparc) 8.60.50242

```

```
Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,??,q]: 1
```

6. Enter **1** to install the **JNPRsbr.pkg**. The system installs the new software.

```
Processing package instance JNPRsbr.pkg
.
.
.
Installation of JNPRsbr.pkg was successful.
```

7. Once the system displays the message **Installation of JNPRsbr.pkg was successful**, repeat this procedure on the next S node. Once the new software is installed on each S node, proceed to [“Configuring the New SBRC Software on the S Node” on page 425](#).

## Configuring the New SBRC Software on the S Node

Configure the new software on each S node, one at a time:

1. Configure the SBRC software on the S node, using the original cluster configuration files.

Execute:

**root@bng-sbrha-1: ./configure**

```
Configuring SBR Software

END USER LICENSE AGREEMENT READ THIS END USER LICENSE AGREEMENT ("AGREEMENT")
.....
.
.
.
.Do you accept the terms in the license agreement? [n]:    y
```

2. Enter **y** to accept the license agreement.

```
-----
SBR 8.60-R1.50765 cluster
on SunOS 5.10 Generic_141444-09 node bng-sbrha-1
```

```
is not configured and processes are down, needs to be configured
-----
```

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.
  2. Generate Cluster Definition  
Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.
  3. Configure Cluster Node  
To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.
  4. Reconfigure RADIUS Server  
Only on SBR nodes, updates the existing SBR configuration.
  5. Create Temporary Cluster  
Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.
  6. Upgrade From Restricted Cluster License  
Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license
- Enter the number of the desired configuration task or quit (2,q): 3

3. Enter **3** to select option **3. Configure Cluster Node**.

```
-----
SBR 8.60-R1.50765 cluster
on SunOS 5.10 Generic_141444-09 node bng-sbrha-1
is not configured and processes are down, needs to be configured
-----

Configuring Cluster Node...
```

4. Specify the name of your existing cluster. In this example the cluster name is "6node\_cluster".

```
Enter SBR cluster name [bng-sbrha-1]: 6node_cluster
```

```
Reading shared configuration from /opt/JNPRshare/install/6node_cluster
.....
```

5. Migrate the configuration from your previous software release to the new software package. In most cases, you simply need to identify your previous installation. In this example, the cluster configuration is named **configure.6node\_cluster.tar** (see the previous step), and is stored in the **/opt/JNPRsbr/radius/install** directory. Once you specify the configuration, SBRC migrates your configuration.

Enter:

```
root@bng-sbrha-1: /opt/JNPRsbr/radius/install
```

```
Reading shared configuration from /opt/JNPRshare/install/6node_cluster
```

```
Enter absolute path to configure.6node_cluster.tar
if you have it, or quit (q) so that you can get it from
/opt/JNPRsbr/radius/install on the first cluster node.
```

```
Generating configuration files
```

```
Reviewing configuration files
```

```
/opt/JNPRsbr/radius/install/tmp/config.ini
```

```
/opt/JNPRsbr/radius/install/tmp/my.cnf
```

```
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
```

```
View (v), accept (a), or reject (r) configuration files: a
```

6. Enter **a** to accept the configuration files.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y
```

7. Enter **y** to continue.

```
Cleaning directories
```

```
/opt/JNPRhadm
```

```
/opt/JNPRmysql
```

```
/opt/JNPRmysqld
```

```
/opt/JNPRndb_mgmd
```

```
/opt/JNPRndbd
```

```

Applying configuration

Initializing Session State Register, please wait a few minutes...
19949 Killed

Do you want to configure Java Runtime Environment for JDBC Feature [n] :

Please enter backup or radius directory from which to migrate.
Enter n for new configuration, s to search, or q to quit
[/opt/JNPRsbr/radius/install/backups/2010:10:21-08:31:21]:

Migrating files, please wait...
-

WARNING: A manual merge may be required between the files located in
/opt/JNPRsbr/radius (working files) and
/opt/JNPRsbr/radius/install/backups/2010:10:21-08:32:18 (shipped files)

```

8. Press **Enter** through the following prompts to keep your configuration.

```

Enter initial admin user (UNIX account must have a valid password) [root]:
Configuring for use with generic database
Do you want to enable "Radius WatchDog" Process? [n]:
Do you want to enable LCI? [n]:
Do you want to configure for use with Oracle? [n]:
Removing oracle references from startup script
Do you want to configure for use with SIGTRAN? [n]:
Removing SIGTRAN references from startup script
Do you want to configure SNMP? [n]:
Configuring Admin GUI Webserver
Compatible Java version found :
Do you want to install custom SSL certificate for Admin WebServer? [n]:

```

9. Enter **e** when prompted to enable the autoboot scripts.

```

Enable (e), disable (d), or preserve (p) autoboot scripts [e]: e

```

10. The main configuration menu is displayed.

Enter **q** to quit.

The SBR Admin Web GUI can be launched using the following URL:  
 https://<servername>:2909

```
-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node bng-sbrha-1(s)
is configured and processes are down, may be reconfigured if desired
-----
```

1. Unconfigure Cluster Node  
     Not used when merely updating existing cluster definitions.
2. Generate Cluster Definition  
     Creates new or updates existing cluster definitions.  
     Modifies the shared directory but does not modify this node.
3. Configure Cluster Node  
     To be preceded by 'Generate Cluster Definition' on one node.  
     Must be invoked on each and every node of the cluster.
4. Reconfigure RADIUS Server  
     Only on SBR nodes, updates the existing SBR configuration.
5. Create Temporary Cluster  
     Used to approximate a cluster using only this one machine.  
     Intended for migration and demonstration purposes only.
6. Upgrade From Restricted Cluster License  
     Used to upgrade from restricted cluster to regular cluster.  
     Removes database restriction on the number of concurrent  
     sessions and enables the addition of an expansion kit license

```
READY: last operation succeeded, node configured.
Enter the number of the desired configuration task or quit (4,q): q
```

## 11. Start the RADIUS process on the S node.

Execute:

```
root@bng-sbrha-1:/opt/JNPRsbr/radius/sbrd start radius
```

```
Starting radius server processes
RADIUS: Process ID of daemon is 14235
RADIUS: Starting DCF system
RADIUS: Configuration checksum: 6A 3B 7A D5
```

```

Document Root: /opt/JNPRsbr/radius/

root@bng-sbrha-1 [/opt/JNPRsbr/radius]
$Starting webserver
. 2016-01-11 20:03:18.493:INFO::main: Logging initialized @3821ms
. . . . . Webserver started
AppleTalk
AppleTalk
AppleTalk
RADIUS:
RADIUS: Steel-Belted Radius licenses
RADIUS: License String:                      Additional Info:
RADIUS: XXXXXXXXXXXXXXXXXXXX                SSR Cluster & Management Expansion
RADIUS: XXXXXXXXXXXXXXXXXXXX                SSR Management Expansion
RADIUS: XXXXXXXXXXXXXXXXXXXX                Run license - lab license - 1k
sessions
RADIUS:
RADIUS: Total concurrent sessions licensed: 1000
RADIUS:
RADIUS: Unlimited Transactions Per Second is Allowed
RADIUS:
RADIUS: Licenses enable this server to:
RADIUS:      Run
RADIUS:      use Multi Node Cluster
RADIUS:      connect '3' management and '2' cluster nodes to database
RADIUS:
RADIUS: Licensed for Steel-Belted Radius Carrier
RADIUS: Attribute Editing enabled
RADIUS: DCF system started

```

12. Check the status of the S node and verify it has properly reconnected to the cluster.

Execute:

**root@bng-sbrha-1: /opt/JNPRsbr/radius/sbrd status**

```

-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node bng-sbrha-1(s)
-----

```

```

-----
Current state of network interfaces:

```

```

10.212.10.65.1646                                Idle

```



```

10.212.10.65.1813          Idle
10.212.10.65.1645          Idle
10.212.10.65.1812          Idle
      *.1812                *. *          0      0 49152      0 LISTEN
      *.1813                *. *          0      0 49152      0 LISTEN
-----

root 14235 radius sbr.xml
root  14305 webserver

```

## Upgrading the Data (D) Nodes with the New Software

After all M, SM, and S nodes in the cluster are upgraded, you need to upgrade the software in D nodes by installing the complete SBR release distribution package or by installing only the latest MySQL package.

### Installing the Complete SBR Release Distribution Package

This section describes the procedure to install the new software on the data nodes by copying the complete SBR release distribution package.

To install the new software on the data node by copying the complete SBR release distribution package:

**NOTE:** You must repeat this procedure on every D node in the cluster, one at a time. This example procedure uses two D nodes, bng-sbr-perfm3000-3 (10.212.10.188) and bng-sbr-perf2 (10.212.10.67).

1. Verify that the `/opt/JNPRmysql` directory contains the directories **install** and **mysql-cluster-advanced-7.4.15-solaris10-sparc-64bit**. If these directories are not present, your installation is using non-default locations for these files. Locate these files and execute the steps in the following procedure in the same directory that contains these files.
2. Log in to the first D node as root, and stop the SSR process. In the following example, the D node is called bng-sbr-perfm3000-3 (10.212.10.188).

Execute:

```
root@bng-sbr-perfm3000-3: ./sbrd stop ssr
```

**Stopping ssr data processes**

3. Log in to a M node and execute `./sbrd status`. Verify that the first D node is disconnected from the other nodes.

Execute:

root@bng-mars: `./sbrd status`

```
-----
SBR 7.60-R1.50242 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node bng-mars.englab.juniper.net(m)
-----

Connected to Management Server at: 10.212.10.68:5235

[ndbd(NDB)] 2 node(s)
id=1 (not connected, accepting connect from 10.212.10.188)
id=2 @10.212.10.67 (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0, *)

[ndb_mgmd(MGM)] 3 node(s)
id=51 @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=52 @10.212.10.68 (mysql-5.6.36 ndb-7.4.15)
id=53 @10.212.10.66 (mysql-5.6.36 ndb-7.4.15)

[mysqld(API)] 5 node(s)
id=61 @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=62 @10.212.10.68 (mysql-5.6.36 ndb-7.4.15)
id=63 @10.212.10.66 (mysql-5.6.36 ndb-7.4.15)
id=100 @10.212.10.219 (mysql-5.6.36 ndb-7.4.15)

hadm      4354 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
hadm      4908 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
hadm      5255 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install

--datadir=/opt/JNPRmysqld/data --plugin-dir=/opt/JNPRmysql/install/lib/plugin
--log-error=/opt/JNPRmysqld/mysqld_safe.err
--pid-file=/opt/JNPRmysqld/mysqld.pid --socket=/opt/JNPRhadm/.mysql.sock
--port=3001
```

4. Log in to the second data node and execute `./sbrd status`, and verify that it is still connected to the cluster. The following is an example of what the status of the second D node should look like:

```
-----
SBR 7.60-R1.50242 cluster 6node_cluster{1s,1sm,2m,2d}
```

```
on SunOS 5.10 Generic_141444-09 node bng-sbr-perf2(d)
-----

root 25604 ndbmttd --initial
--ndb-connectstring=nodeid=2;10.212.10.213:5235,10.212.10.68:5
root 25605 ndbmttd --initial
--ndb-connectstring=nodeid=2;10.212.10.213:5235,10.212.10.68:5
```

5. Log in to the first data node and then move your original installation to a different location to prevent the removal of files. Be sure not to move the **/opt/JNPRshare** directory, which should contain your original cluster configuration files. For example, if you previously installed Release 7.6.0, execute the following as root:

- a. Navigate to the opt directory: **root@bng-sbr-perfm3000-3: cd /opt.**
- b. Create a new directory: **root@bng-sbr-perfm3000-3: mkdir 760.**
- c. Move the configuration files to the new directory: **root@bng-sbr-perfm3000-3: mv JNPRsbr JNPRmysql JNPRmysqld JNPRhadm JNPRndb\_mgmd JNPRndbd 760.**
- d. Remove the old software package: **root@bng-sbr-perfm3000-3: pkgrm JNPRsbr.**

6. Navigate to the **/opt/tmp** directory and install the new SBRC software package.

Execute:

**root@bng-sbr-perfm3000-3: pkgadd -d.**

```
The following packages are available:
  1  JNPRsbr.pkg      JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier
Cluster Edition)
                                (sparc) 8.60.50242

Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,?,q]: 1
```

7. Enter **1** to select the **JNPRsbr.pkg**.  
The script resumes.

```
Processing package instance <JNPRsbr.pkg> from </tmp>
```

8. Confirm the installation directory.

Depending on the system configuration, the script prompts you to create the **/opt/JNPRsbr** directory if it does not exist, overwrite an already-extracted package, or any of several other prompts.

```
The selected base directory </opt/JNPRsbr> must exist before installation is
attempted.
```

```
Do you want this directory created now [y,n,?,q]
```

Answer the prompt appropriately (or change the extraction path if necessary) so that the script can proceed.

To accept the default directory as a target, enter **y**.

The script resumes.

```
Using </opt/JNPRsbr> as the package base directory.
#Processing package information.
#Processing system information.
    48 package pathnames are already properly installed.
#Verifying disk space requirements.
#Checking for conflicts with packages already installed.
#Checking for setuid/setgid programs.
```

```
This package contains scripts which will be executed with super-user
permission during the process of installing this package.
```

```
Do you want to continue with the installation of <JNPRsbr> [y,n,?]
```

9. Enter **y** to confirm that you want to continue to install the package.

```
Installing JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster
Edition) as <JNPRsbr>
```

```
## Executing preinstall script.
## Installing part 1 of 1.
.
.
.
[ verifying class <none> ]
## Executing postinstall script.
```

```
Newly installed server directory will be backed up as:
/opt/JNPRsbr/radius/install/backups/2009:03:31-00:34:06
```

```
Installation of <JNPRsbr> was successful.
```

Once the system displays the message **Installation of <JNPRsbr> was successful**, the installation of the new software is complete on the first data node in the cluster. Once the software installation is completed, configure the new software on the data node and start the SSR process in the data node before repeating the installation steps on the other data nodes in the cluster. For more information about the software configuration, see [“Configuring the New Software on the Data Node” on page 439](#).

### Installing Only the Latest MySQL Package

To install the new software on the data node by copying only the latest MySQL package:



**CAUTION:** During your maintenance window, execute the following steps. If you run out of time during the maintenance window, it is possible to upgrade only some of the data nodes. This partially upgraded cluster has not been extensively tested by Juniper Networks and is not supported, so every attempt must be made to upgrade the entire cluster. It is recommended to install the new software by copying the complete SBR release distribution package (see [“Installing the Complete SBR Release Distribution Package” on page 431](#)).

**NOTE:** You must repeat this procedure on every D node in the cluster, one at a time.

1. Verify that the `/opt/JNPRmysql` directory contains the directories **install** and **mysql-cluster-advanced-7.4.15-solaris10-sparc-64bit**. If these directories are not present, your installation is using non-default locations for these files. Locate these files and execute the steps in the following procedure in the same directory that contains these files.
2. Verify that `/opt/JNPRsbr/radius` contains the **sbrd** file. If this file is not present, your installation is using a non-default location for the SBRC software distribution package. You need to locate the **sbrd** file and adjust the instructions below with the correct path to the file.
3. Change to the directory located in Step 1, (usually, `/opt/JNPRmysql`) `cd /opt/JNPRmysql`.
4. Unzip and un-tar the new MySQL package:

Execute:

```
gunzip -c /opt/tmp/JNPRsbr.pkg/reloc/radius/install/ndb/bin/*mysqlcluster*.tar.gz \ | tar -xf -
```

5. Log in to the first D node as root, and stop the SSR process. In the following example, the D node is called bng-sbr-perfm3000-3.

Execute:

```
root@bng-sbr-perfm3000-3: ./sbrd stop ssr
```

### Stopping ssr data processes

6. Check the status of the D node.

Execute:

```
root@bng-sbr-perfm3000-3: ./sbrd status
```

```
-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node bng-sbr-perfm3000-3(d)
-----
```

7. Log in to the SM node and execute **./sbrd status**, and verify that the D node is disconnected from the other nodes. The following is an example of what the status should look like:

```
-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node sbr1.englab.juniper.net(sm)
-----

Connected to Management Server at: 10.212.10.213:5235

[ndbd(NDB)]      2 node(s)
id=1   (not connected, accepting connect from 10.212.10.188)
id=2   @10.212.10.67 (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0, Master)

[ndb_mgmd(MGM)]  3 node(s)
id=51 @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=52 @10.212.10.68 (mysql-5.6.36 ndb-7.4.15)
id=53 @10.212.10.66 (mysql-5.6.36 ndb-7.4.15)
[mysqld(API)]   5 node(s)
id=61 @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=62 @10.212.10.68 (mysql-5.6.36 ndb-7.4.15)
id=63 @10.212.10.66 (mysql-5.6.36 ndb-7.4.15)
id=100 @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=101 @10.212.10.65 (mysql-5.6.36 ndb-7.4.15)
```

```

10.212.10.213.1646                               Idle
10.212.10.213.1813                               Idle
10.212.10.213.1645                               Idle
10.212.10.213.1812                               Idle
      *.1812      *.*      0      0 49152      0 LISTEN
      *.1813      *.*      0      0 49152      0 LISTEN

      hadm  4945 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
--datadir=/o
      hadm  4776 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
      hadm  4755 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
      root  8422 radius sbr.xml

```

8. Remove the existing symbolic link and create a new symbolic link for the newly extracted MySQL package. For example:

Execute:

**rm install**

**ln -s /opt/760/mysql-cluster-advanced-7.4.15-solaris10-sparc-64bit install**

9. Start the SSR process on the data node.

Execute:

**root@bng-sbr-perfm3000-3: /opt/JNPRsbr/radius/sbrd start ssr**

10. Check the status of the upgraded data node.

Execute:

**root@bng-sbr-perfm3000-3: ./sbrd status**

```

-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node bng-sbr-perfm3000-3(d)
-----

```

```

-----
Current state of network interfaces:
-----

```

```

      root  8779 ndbmttd --initial
--ndb-connectstring=nodeid=1;10.212.10.213:5235,10.212.10.68:5
      root  8778 ndbmttd --initial
--ndb-connectstring=nodeid=1;10.212.10.213:5235,10.212.10.68:5

```

Examine the display. The data node remains in a 'starting' state until it has copied all the data from the primary node with which it is paired. This process may take 30 minutes for a typical load. When the data node has completed, it will display as being reconnected to the cluster, as in the above example.



**CAUTION:** Once the data node finishes 'starting' and is shown in the same state as the other data nodes in the cluster, the upgrade of this data node is complete and the cluster is once again in a highly available state. It is imperative not to begin upgrading the next data node until this node has re-synchronized in this manner. Otherwise, it is possible to lose the cluster database contents and experience an outage.

11. While waiting for the data node to complete the starting process and synchronize, monitor the connection state of the SM nodes in the cluster by repeating the **./sbrd status** command. You should see a similar display, as follows:

```
-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node sbr1.englab.juniper.net(sm)
-----

Connected to Management Server at: 10.212.10.213:5235

[ndbd(NDB)]      2 node(s)
id=1      @10.212.10.188  (mysql-5.6.36 ndb-7.4.15, starting, Nodegroup: 0)
id=2      @10.212.10.67  (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0, *)

[ndb_mgmd(MGM)]  3 node(s)
id=51     @10.212.10.213  (mysql-5.6.36 ndb-7.4.15)
id=52     @10.212.10.68  (mysql-5.6.36 ndb-7.4.15)
id=53     @10.212.10.66  (mysql-5.6.36 ndb-7.4.15)

[mysqld(API)]    5 node(s)
id=61     @10.212.10.213  (mysql-5.6.36 ndb-7.4.15)
id=62     @10.212.10.68  (mysql-5.6.36 ndb-7.4.15)
id=63     @10.212.10.66  (mysql-5.6.36 ndb-7.4.15)
id=100    @10.212.10.213  (mysql-5.6.36 ndb-7.4.15)
id=101    @10.212.10.65  (mysql-5.6.36 ndb-7.4.15)

-----

Current state of network interfaces:
```



```

10.212.10.213.1646      Idle
10.212.10.213.1813      Idle
10.212.10.213.1645      Idle
10.212.10.213.1812      Idle
    *.1812              *. *          0      0 49152      0 LISTEN
    *.1813              *. *          0      0 49152      0 LISTEN
-----

    hadm  4761 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
    hadm  4779 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
    hadm  4968 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
--datadir=/o
    root  5405 radius sbr.xml
    root  5475 webserver

```

If the SM nodes suddenly disconnect, this is a symptom of a serious error in the SSR cluster that can occur under rare circumstances. If this occurs, you must manually restart the affected node by executing **/opt/JNPRsbr/radius/sbrd restart radius** to avoid an outage.

12.Repeat this upgrade procedure until each data node has been upgraded.

**NOTE:** Software configuration is not required if you upgrade the software in D node by installing the MySQL package.

## Configuring the New Software on the Data Node

Configure the new SBRC software on the D nodes by running the configuration script.

**NOTE:** You must repeat this procedure on every data node in the cluster.

1. As root, navigate to the directory where you installed the Steel-Belted Radius Carrier package.

For example, navigate to the **radius/install** subdirectory.

Execute:

```
root@bng-sbr-perfm3000-3: cd /opt/JNPRsbr/radius/install/
```

2. Execute the configuration script to configure the SBRC software.

Execute:

```
root@bng-sbr-perfm3000-3: ./configure
```

### 3. Review and accept the Steel-Belted Radius Carrier license agreement.

Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

```
-----
SBR 8.60-R1.50765 cluster
on SunOS 5.10 Generic_141444-09 node bng-sbr-perfm3000-3
is not configured and processes are down, needs to be configured
-----

1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.

2.  Generate Cluster Definition
    Creates new or updates existing cluster definitions.
    Modifies the shared directory but does not modify this node.

3.  Configure Cluster Node
    To be preceded by 'Generate Cluster Definition' on one node.
    Must be invoked on each and every node of the cluster.

4.  Reconfigure RADIUS Server
    Only on SBR nodes, updates the existing SBR configuration.

5.  Create Temporary Cluster
    Used to approximate a cluster using only this one machine.
    Intended for migration and demonstration purposes only.

6.  Upgrade From Restricted Cluster License
    Used to upgrade from restricted cluster to regular cluster.
    Removes database restriction on the number of concurrent
    sessions and enables the addition of an expansion kit license
```

### 4. From the menu of configuration tasks, enter **3** to specify **Configure Cluster Node**.

```
-----
SBR 8.60-R1.50765 cluster
on SunOS 5.10 Generic_141444-09 node bng-sbr-perfm3000-3
is not configured and processes are down, needs to be configured
-----
```

```
Configuring Cluster Node...
```

5. When prompted, enter the name of your existing cluster. In this example the cluster name is "6node\_cluster".

```
Enter SBR cluster name [sbrha]: 6node_cluster
```

You are prompted to verify whether you want to proceed, unless the script detects any unusual installation conditions (a pre-existing directory, for example). In some cases, you may be prompted to resolve or ignore them.

6. The system reads the configuration files that you copied to the server and prompts you to change some settings to adapt them to this server.

Enter **a** to accept the files.

```
Reading shared configuration from /opt/JNPRshare/install/6node_cluster

Generating configuration files

Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
View (v), accept (a), or reject (r) configuration files: a
```

7. The system displays the following warning and prompts you to continue. Enter **y** to continue.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? [n]: y
```

```
Cleaning directories
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRndbd

Applying configuration
```

Initializing Session State Register, please wait a few minutes...

**NOTE:** Expect package is required to be installed on RHEL 7.3 or later (see [“Linux” on page 47](#) for supported RHEL versions). SBR Carrier has been tested only with the Expect-5.45 version. Expect package is installed on Solaris 11.3.36.10.0 or later version by default (see [“Solaris” on page 47](#) for supported Solaris versions).

Refer the Release Notes document for more information on the Solaris 11.4.25.0.1.75.3 version.

## 8. When SSR is finished, you are returned to the main configuration menu.

```
-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node bng-sbr-perfm3000-3(d)
is configured and processes are down, may be reconfigured if desired
-----
```

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.
2. Generate Cluster Definition  
Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.
3. Configure Cluster Node  
To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.
4. Reconfigure RADIUS Server  
Only on SBR nodes, updates the existing SBR configuration.
5. Create Temporary Cluster  
Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.
6. Upgrade From Restricted Cluster License  
Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

```
READY: last operation succeeded, node configured.
Enter the number of the desired configuration task or quit (2,q): q
```

Enter **q** to quit.

9. Start the SSR process on the data node.

Execute:

```
root@bng-sbr-perfm3000-3: /opt/JNPRsbr/radius/sbrd start ssr
```

10. Check the status of the upgraded data node.

Execute:

```
root@bng-sbr-perfm3000-3: ./sbrd status
```

```
-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node bng-sbr-perfm3000-3(d)
-----

-----
Current state of network interfaces:
-----

root 8779 ndbmttd --initial
--ndb-connectstring=nodeid=1;10.212.10.213:5235,10.212.10.68:5
root 8778 ndbmttd --initial
--ndb-connectstring=nodeid=1;10.212.10.213:5235,10.212.10.68:5
```

Examine the display. The data node remains in a 'starting' state until it has copied all the data from the primary node with which it is paired. This process may take 30 minutes for a typical load. When the data node has completed, it will display as being reconnected to the cluster, as in the above example.



**CAUTION:** Once the data node finishes 'starting' and is shown in the same state as the other data nodes in the cluster, the upgrade of this data node is complete and the cluster is once again in a highly available state. It is imperative not to begin upgrading the next data node until this node has re-synchronized in this manner. Otherwise, it is possible to lose the cluster database contents and experience an outage.

11. While waiting for the data node to complete the starting process and synchronize, monitor the connection state of the SM nodes in the cluster by repeating the `./sbrd status` command. You should see a similar display, as follows:

```
-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node sbr1.englab.juniper.net(sm)
-----
```

```
Connected to Management Server at: 10.212.10.213:5235
```

```
[ndbd(NDB)]      2 node(s)
id=1    @10.212.10.188  (mysql-5.6.36 ndb-7.4.15, starting, Nodegroup: 0)
id=2    @10.212.10.67  (mysql-5.5.20 ndb-7.2.5, Nodegroup: 0, *)
```

```
[ndb_mgmd(MGM)] 3 node(s)
id=51   @10.212.10.213  (mysql-5.6.36 ndb-7.4.15)
id=52   @10.212.10.68  (mysql-5.6.36 ndb-7.4.15)
id=53   @10.212.10.66  (mysql-5.6.36 ndb-7.4.15)
```

```
[mysqld(API)]   5 node(s)
id=61    @10.212.10.213  (mysql-5.6.36 ndb-7.4.15)
id=62    @10.212.10.68  (mysql-5.6.36 ndb-7.4.15)
id=63    @10.212.10.66  (mysql-5.6.36 ndb-7.4.15)
id=100   @10.212.10.213  (mysql-5.6.36 ndb-7.4.15)
id=101   @10.212.10.65  (mysql-5.6.36 ndb-7.4.15)
```

```
-----
Current state of network interfaces:
```

```
10.212.10.213.1646      Idle
10.212.10.213.1813      Idle
10.212.10.213.1645      Idle
10.212.10.213.1812      Idle
      *.1812             *.*             0         0 49152      0 LISTEN
      *.1813             *.*             0         0 49152      0 LISTEN
-----
```

```
hadm 4761 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
hadm 4779 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
hadm 4968 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
--datadir=/o
```

```

root 5405 radius sbr.xml
root 5475 webserver

```

If the SM nodes suddenly disconnect, this is a symptom of a serious error in the SSR cluster that can occur under rare circumstances. If this occurs, you must manually restart the affected node by executing **/opt/JNPRsbr/radius/sbrd restart radius** to avoid an outage.

12. Repeat this upgrade procedure until each data node has been upgraded; see [“Upgrading the Data \(D\) Nodes with the New Software” on page 431](#).

After all nodes are upgraded, you should see a similar output as follows:

```

-----
SBR 8.60-R1.50765 cluster 6node_cluster{1s,1sm,2m,2d}
on SunOS 5.10 Generic_141444-09 node sbr1.englab.juniper.net(sm)
-----

Connected to Management Server at: 10.212.10.213:5235

[ndbd(NDB)]      2 node(s)
id=1   @10.212.10.188 (mysql-5.6.36 ndb-7.4.15, Nodegroup: 0, *)
id=2   @10.212.10.67  (mysql-5.6.36 ndb-7.4.15, Nodegroup: 0)

[ndb_mgmd(MGM)]  3 node(s)
id=51  @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=52  @10.212.10.68  (mysql-5.6.36 ndb-7.4.15)
id=53  @10.212.10.66  (mysql-5.6.36 ndb-7.4.15)

[mysqld(API)]    5 node(s)
id=61  @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=62  @10.212.10.68  (mysql-5.6.36 ndb-7.4.15)
id=63  @10.212.10.66  (mysql-5.6.36 ndb-7.4.15)
id=100 @10.212.10.213 (mysql-5.6.36 ndb-7.4.15)
id=101 @10.212.10.65  (mysql-5.6.36 ndb-7.4.15)

-----

Current state of network interfaces:

10.212.10.213.1646      Idle
10.212.10.213.1813      Idle
10.212.10.213.1645      Idle
10.212.10.213.1812      Idle
      *.1812             *. *             0      0 49152      0 LISTEN

```

```

      *.1813          *. *          0          0 49152          0 LISTEN
-----

hadm  4761 ndb_mgmd --config-cache=0 --configdir=/opt/JNPRhadm
hadm  4779 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
hadm  4968 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
--datadir=/o
root  5405 radius sbr.xml
root  5475 webserver

```

## Launching Web GUI

Now that the RADIUS process is running, you can complete the configuration using Web GUI. For details, see [“Basic SBR Carrier Node Configuration” on page 138](#). For complete details, see the *SBR Carrier Administration and Configuration Guide*.



# Upgrading Your Cluster Using the Backup, Destroy, and Re-Create Method

## IN THIS CHAPTER

- Overview of the Backup, Destroy, and Re-Create Upgrade Method | 448
- Capturing Your Current Cluster Configuration | 448
- Backing Up the State of the Sessions in Your SSR Database | 450
- Destroying the Cluster Database | 450
- Stopping All Processes on All Nodes | 451
- Installing the New SBRC Software on the Data Nodes | 452
- Installing the New SBRC Software on the SM Nodes | 455
- Configuring the First SM Node and Creating a New Cluster Definition | 457
- Configuring the SBRC Software on the SM Nodes | 463
- Configuring the SBRC Software on the Data Nodes | 469
- Starting the SSR Processes | 473
- Creating and Restoring the SSR Database | 476
- Starting the RADIUS Process | 478
- Launching Web GUI | 480

This section describes the backup, destroy, and re-create upgrade method, which you can use to upgrade servers running in a SSR cluster. For information about upgrading servers running standalone, see [“Installing and Configuring a SBR Carrier Standalone Server”](#) on page 63.

These topics are included in this section:

## Overview of the Backup, Destroy, and Re-Creat Upgrade Method

**NOTE:** This upgrade procedure requires downtime of the cluster, which can be mitigated by temporarily switching your cluster traffic to a single transition server. For details on using a transition server, see [“Using a Transition Server to Mitigate Downtime While Upgrading Your Cluster” on page 369](#). Although the use of a transition server can mitigate downtime, you will experience slower performance since a single server does not provide the same performance as a cluster.



**CAUTION:** This procedure requires you to back up the cluster database, destroy the database, install and configure the new SBRC software, and re-create your database. You should schedule a maintenance window when performing this upgrade, even if you use a transition server to mitigate a complete outage. If you are using a transition server, do not proceed with this upgrade until you have tested the operation and performance of the transition server.

This upgrade method results in the data nodes using a multi-threaded process, which provides improved performance over the single-threaded data node processes in previous SBR Carrier software releases.

The example procedures in this section assume your cluster consists of two SM nodes and two data (D) nodes as follows:

- First SM node, uranus-js (10.13.20.89)
- Second SM node, sbrha-3 (10.13.20.77)
- First D node, sbrha-9 (10.13.20.83)
- Second D node, sbrha-7 (10.13.20.81)

## Capturing Your Current Cluster Configuration

To prepare for the upgrade, capture your current cluster configuration:

1. (Optional) If you previously upgraded the D nodes in your cluster to Release 8.6.0 by pushing the package only as described in Step 3 of [“Preparation” on page 399](#), you need to copy the **sbrd** script from one SM node to each data node. If you upgraded using the **pkgadd** command, you do not need to perform this step.

2. Log in to each node in your cluster as root, and verify that it is operating correctly. Print out the status of each node in the cluster. For example:

Execute:

**root@uranus-js: /opt/JNPRsbr/radius./sbrd status**

```
-----
SBR 8.60.50006 cluster sbrqa{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net(sm)
-----

Connected to Management Server at: 10.13.20.89:5235

[ndbd(NDB)]      2 node(s)
id=1    @10.13.20.83  (mysql-5.6.36 ndb-7.4.15, Nodegroup: 0)
id=2    @10.13.20.81  (mysql-5.6.36 ndb-7.4.15, Nodegroup: 0, Master)

[ndb_mgmd(MGM)]  2 node(s)
id=51   @10.13.20.89  (mysql-5.6.36 ndb-7.4.15)
id=52   @10.13.20.77  (mysql-5.6.36 ndb-7.4.15)

[mysqld(API)]    4 node(s)
id=61   @10.13.20.89  (mysql-5.6.36 ndb-7.4.15)
id=62   @10.13.20.77  (mysql-5.6.36 ndb-7.4.15)
id=100  @10.13.20.89  (mysql-5.6.36 ndb-7.4.15)
id=101  @10.13.20.77  (mysql-5.6.36 ndb-7.4.15)

10.13.20.89.1646          Idle
10.13.20.89.1813          Idle
10.13.20.89.1645          Idle
10.13.20.89.1812          Idle
      *.1812              *.*              0      0 49152      0 LISTEN
      *.1813              *.*              0      0 49152      0 LISTEN

  hadm 25069 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
--datadir=/o
  hadm 25004 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
  hadm 24983 ndb_mgmd
  root 28131 radius sbr.xml
```

## Backing Up the State of the Sessions in Your SSR Database

To back up the state of the sessions in your SSR database:

- Log in to one of the SM nodes as **hadm**, make a backup of your cluster database, and copy it to the **opt/tmp** directory. For example, to log in to the SM node called “**uranus-js**” and make a backup of the database:

Execute:

```
hadm@uranus-js: mysqldump -p SteelBeltedRadius > opt/tmp/backup_sbr.sql
```

## Destroying the Cluster Database

Destroy the existing cluster database:

1. Log in to each SM node as **root**, and stop the **RADIUS** process.

Execute:

```
root@uranus-js: ./sbrd stop radius
```

```
Stopping radius server processes
waiting for radius
10 seconds elapsed, still waiting
radius stopped
```

2. From one of the SM nodes, log in as **hadm** and destroy the cluster database.

Execute:

```
hadm@uranus-js: ./DestroyDB.sh
```

```
SBRs must be offline; OK? <yes|no>  yes
This will destroy the "SteelBeltedRadius" database; OK? <yes|no>  yes
Really? <yes|no>  yes
Database "SteelBeltedRadius" destroyed.
hadm@uranus-js:
```

## Stopping All Processes on All Nodes

On each SM node, stop both the SSR and RADIUS process. On each D node stop, the SSR process.

1. Stop the SSR process and RADIUS process on the SM node.

- Execute:

```
root@uranus-js: /opt/JNPRsbr/radius]$ ./sbrd stop ssr
```

```
Stopping ssr auxiliary processes
Stopping ssr management processes
Connected to Management Server at: 10.13.20.89:5235
Node 51 has shutdown.
Disconnecting to allow Management Server to shutdown

Killing remaining ssr processes
    hadm 24983 ndb_mgmd
```

- Shut down the RADIUS process on the SM node.

Execute:

```
root@uranus-js: /opt/JNPRsbr/radius]$ ./sbrd stop radius
```

```
waiting for radius
10 seconds elapsed, still waiting
radius stopped
```

- Repeat for each SM node.

2. Stop the SSR process on each data node.

- Log in to the data node and execute:

```
sbrha-7: /opt/JNPRsbr/radius]$ ./sbrd stop ssr
```

```
Stopping ssr data processes
```

- Repeat for each Data node.

## Installing the New SBRC Software on the Data Nodes

**NOTE:** You only need to perform the steps in this section if you did not previously perform the rolling restart upgrade procedure (see [“Upgrading Your Cluster Using the Rolling Restart Method” on page 396](#)).

You need to uninstall and remove the current software package and install new software package on *each* D node (one at a time) and M-only node (not applicable in this example procedure).

### Uninstalling and Removing the Software Package on the Data Nodes

You must perform this procedure on each data node, one at a time.

1. To uninstall the current software package, navigate to the `/opt/JNPRsbr/radius/install` directory.

Execute:

**root@sbrha-9: ./unconfigure**

```
Unconfiguring SBR Software

-----
SBR 8.60.50006 cluster sbrqa{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-9.englab.juniper.net(d)
is configured and processes are down, may be reconfigured if desired
-----

Unconfiguring Cluster Node...

WARNING: You are about to unconfigure this node.
Are you sure that you wish to continue? (y,n): y

Cleaning directories
/opt/JNPRhadm
/opt/JNPRmysql
/opt/JNPRndbd

Locating shared directory...
drwxrwxr-x  2 hadm  hadmg      512 Oct 14 17:02 /opt/JNPRshare/install/sbrqa

WARNING: If you remove the shared directory for this cluster, you
will either have to recover the data from another cluster node or
```

```

reconfigure the entire cluster again. This is neither necessary
nor recommended if you are updating an existing configuration.
Remove the shared directory for this cluster? [n]:

Locating OS user account and home directory...
hadm:x:54321:12345:SBR Software:/opt/JNPRhadm:/bin/bash
hadmg::12345:
drwxrwx---  2 hadm      hadmg          512 Oct 21 19:19 /opt/JNPRhadm

WARNING: If you remove the OS user account hadm you will have
to recreate it, the associated OS group account hadmg, and
the associated home directory /opt/JNPRhadm
This is neither necessary nor recommended if you are updating
an existing configuration. Remove the OS user account? [n]:

Unconfigured

```

## 2. Remove the software package.

Execute:

**root@sbrha-9: pkgrm JNPRsbr**

```

The following package is currently installed:
  JNPRsbr  JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster
Edition)
          (sparc) 8.60.50006

Do you want to remove this package? [y,n,?,q] y

## Removing installed package instance JNPRsbr

This package contains scripts which will be executed with super-user
permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] y
## Verifying package JNPRsbr dependencies in global zone
## Processing package information.
## Executing preremove script.
Existing server directory will be backed up as:

```

```

/opt/JNPRsbr/radius/install/backups/2010:10:21-19:21:20
## Removing pathnames in class none
/opt/JNPRsbr/radius/install non-empty directory not removed
/opt/JNPRsbr/radius non-empty directory not removed
## Executing postremove script.
## Updating system information.

Removal of JNPRsbr was successful.

```

3. Repeat Step 1 and Step 2 on the next data node.

## Installing the New Software Package on the Data Nodes

You must install the new software on each data node, one at a time.

- To install the new software package, navigate to the folder where you unzipped it, for example the **/opt/tmp** directory.

Execute:

**root@sbrha-9: pkgadd -d.**

```

The following packages are available:
  1  JNPRsbr.pkg      JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier
Cluster Edition)
                                (sparc) 8.60.50006

Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,??,q]: 1

Processing package instance JNPRsbr.pkg from /opt/tmp/gdir

JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster Edition)(sparc)
8.60.50006
(C) Copyright 1996-2016 Juniper Networks, Inc. See license.txt
Using /opt/JNPRsbr as the package base directory.
## Processing package information.
## Processing system information.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.

The following files are already installed on the system and are being
used by another package:
* /opt/JNPRsbr/radius attribute change only

```



```

* /opt/JNPRsbr/radius/install attribute change only

* - conflict with a file which does not belong to any package.

Do you want to install these conflicting files [y,n,?,q] y
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of JNPRsbr [y,n,?] y

Installing JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster Edition)
as JNPRsbr

## Executing preinstall script.
## Installing part 1 of 1.
/opt/JNPRsbr/radius/3GPP.dct
/opt/JNPRsbr/radius/3GPP2.dct
/opt/JNPRsbr/radius/3comsw.dct
/opt/JNPRsbr/radius/3gpp.ini
/opt/JNPRsbr/radius/3gpp2.ini
.
.
.

/opt/JNPRsbr/radius/wimax.ini
/opt/JNPRsbr/radius/wimaxAttributeProcessor.ctrl
/opt/JNPRsbr/radius/wimaxAttributeProcessor.so
/opt/JNPRsbr/radius/xylan.dct
[ verifying class none ]
## Executing postinstall script.
Newly installed server directory will be backed up as:
/opt/JNPRsbr/radius/install/backups/2010:10:21-19:24:45

Installation of JNPRsbr was successful.

```

## Installing the New SBRC Software on the SM Nodes

Install the new software on each SM nodes.

1. Install the new software package on each SM node, one at a time:

Execute:

**root@uranus-js: pkgadd -d.**

```
The following packages are available:
  1  JNPRsbr.pkg      JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier
Cluster Edition)

                                (sparc) 8.60.50006

Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,??,q]: 1

Processing package instance  JNPRsbr.pkg  from  /opt/gdir

JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster Edition)(sparc)
8.60.50006
(C) Copyright 1996-2016 Juniper Networks, Inc.  See license.txt
Using /opt/JNPRsbr  as the package base directory.
## Processing package information.
## Processing system information.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.

The following files are already installed on the system and are being
used by another package:
* /opt/JNPRsbr/radius  attribute change only
* /opt/JNPRsbr/radius/install  attribute change only

* - conflict with a file which does not belong to any package.

Do you want to install these conflicting files [y,n,?,q] y
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of  JNPRsbr  [y,n,?] y

Installing JNPRsbr - Juniper Networks Steel-Belted Radius (Carrier Cluster
Edition) as  JNPRsbr

## Executing preinstall script.
## Installing part 1 of 1.
```

```

/opt/JNPRsbr/radius/3GPP.dct
/opt/JNPRsbr/radius/3GPP2.dct
/opt/JNPRsbr/radius/3comsw.dct
/opt/JNPRsbr/radius/3gpp.ini
.
.
.

/opt/JNPRsbr/radius/wimax.ini
/opt/JNPRsbr/radius/wimaxAttributeProcessor.ctrl
/opt/JNPRsbr/radius/wimaxAttributeProcessor.so
/opt/JNPRsbr/radius/xylan.dct
[ verifying class none ]
## Executing postinstall script.
Newly installed server directory will be backed up as:
/opt/JNPRsbr/radius/install/backups/2010:10:22-06:27:45

Installation of JNPRsbr was successful.

```

2. Repeat on the next SM node.

## Configuring the First SM Node and Creating a New Cluster Definition

To configure the new software on the first SM node and create a new cluster definition:

1. Configure the new software with the new cluster definition on each SM node.

Execute:

```
root@uranus-js: ./configure
```

2. Review and accept the Steel-Belted Radius Carrier license agreement and create a new cluster definition by entering **2 Generate Cluster Definition**.

Press the spacebar to move from one page to the next. When you are prompted to accept the terms of the license agreement, enter **y**.

**Do you accept the terms in the license agreement? [n] y**

```

-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net
is not configured and processes are down, needs to be configured
-----

```

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.
  2. Generate Cluster Definition  
Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.
  3. Configure Cluster Node  
To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.
  4. Reconfigure RADIUS Server  
Only on SBR nodes, updates the existing SBR configuration.
  5. Create Temporary Cluster  
Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.
  6. Upgrade From Restricted Cluster License  
Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license
- Enter the number of the desired configuration task or quit (2,q): 2

3. Create a new cluster definition by entering **2 Generate Cluster Definition**. This example creates a cluster named “junCluster” that contains two SM nodes named “uranus-js” and “sbrha-3”, and two data nodes named “sbrha-9” and “sbrha-7”.

**NOTE:** If you use the original cluster name, the system attempts to use the original configuration files, which may contain pre-8.6.0 values. You must either back up and remove the old cluster definition, or use a new cluster name. However, you can use the same node names.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net
is not configured and processes are down, needs to be configured
-----
```

```
Generating Cluster Definition...
```

When prompted, enter the new name for the cluster and the node names.

```
Enter SBR cluster name [uranus-js]: junCluster
```

The SBR Cluster Starter Kit license allows you to create a minimal cluster of 2 SBR nodes, 2 management nodes, and 2 data nodes. When each node is installed on a separate machine the cluster topology is denoted as {2s,2m,2d}. When SBR nodes are paired with management nodes on the same machines the cluster topology is denoted as {2sm,2d}.

An optional SBR Cluster Management Expansion Kit allows you to add a third management node for {2sm,1m,2d} and an optional Data Expansion Kit allows you to add 2 more data nodes for {2sm,1m,4d} clusters. Additional SBR licenses allow you to add up to 18 more SBR nodes to obtain a maximal cluster {18s,2sm,1m,4d} and/or enable extra features.

While it is not difficult to add management and/or SBR nodes to an existing cluster, adding data nodes is more difficult and may require you to shutdown the entire cluster as opposed to a rolling restart.

4. When prompted, enter the license key.

```
Enter Starter Kit license: XXXX XXXX XXXX XXXX
```

```
Enter Management Expansion Kit license, if any:
```

```
Enter Data Expansion Kit license, if any:
```

5. When prompted, enter the number of SBR (S) nodes.

```
Enter total number of SBR nodes to be configured [2]:2
```

6. When prompted, enter the number of management (M) nodes.

```
Enter number of management nodes to be paired with SBR nodes [2]:2
```

```
Creating cluster junCluster{0s,2sm,0m,2d}
will require 4 machines total. Do you wish to continue? [y]:
```

7. When prompted whether you wish to continue, enter **y** to continue.

```
All cluster nodes will share the same Session State Register (SSR).
Do you wish to continue? [y]:y
```

8. When prompted whether you wish to continue, enter **y** to continue.

```
All cluster nodes will share the same Session State Register (SSR).
Setting password for SSR admin account hadmsql
Password:
Again:
Setting password for SSR software account hadmsbr
Password:
Again:
```

9. When prompted, enter the password for the account.

```
Information will now be gathered for each machine in the cluster.
You will have a chance to review all information at least once
before any machines are modified.
```

10. Enter the information to re-create each node in your cluster.

```
-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net
Partial configuration at present is {0s,0sm,0m,0d} of {0s,2sm,0m,2d}
-----
IMPORTANT: node names must be entered as reported by 'uname -n'.
```

```

Enter node name [junCluster-1]: uranus-js.englab.juniper.net
Enter node type (sm,d) [sm]:
Enter SBR node ID (100-149) [100]:
Enter SBR node IP address by which it is known to management nodes.
Enter SBR node IP address: 10.13.20.89
Enter SBR licenses meant only for this particular SBR node.
Enter one license per line and an empty line when finished.
Enter SBR full license: XXXX XXXX XXXX XXXX
Enter SBR feature license:
Enter MGMT node ID (51-59) [51]:
Enter MGMT node IP address by which it is known to other nodes.
Enter MGMT node IP address: 10.13.20.89

```

```

-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net
Partial configuration at present is {0s,1sm,0m,0d} of {0s,2sm,0m,2d}
-----

```

```

IMPORTANT: node names must be entered as reported by 'uname -n'.
Enter node name [junCluster-2]: sbrha-3.englab.juniper.net
Enter node type (sm,d) [sm]:
Enter SBR node ID (100-149) [101]:
Enter SBR node IP address by which it is known to management nodes.
Enter SBR node IP address: 10.13.20.77
Enter SBR licenses meant only for this particular SBR node.
Enter one license per line and an empty line when finished.
Enter SBR full license: XXXX XXXX XXXX XXXX
Enter SBR feature license:
Enter MGMT node ID (51-59) [52]:
Enter MGMT node IP address by which it is known to other nodes.
Enter MGMT node IP address: 10.13.20.77

```

```

-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net
Partial configuration at present is {0s,2sm,0m,0d} of {0s,2sm,0m,2d}
-----

```

```

IMPORTANT: node names must be entered as reported by 'uname -n'.
Enter node name [junCluster-3]: sbrha-9.englab.juniper.net
Enter node type (d) [d]:
Enter DATA node ID (1-40) [1]:
Enter DATA node IP address by which it is known to management nodes.
Enter DATA node IP address: 10.13.20.83

```

```

-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}

```

```
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net
Partial configuration at present is {0s,2sm,0m,1d} of {0s,2sm,0m,2d}
```

```
-----
IMPORTANT: node names must be entered as reported by 'uname -n'.
Enter node name [junCluster-4]: sbrha-7.englab.juniper.net
Enter node type (d) [d]:
Enter DATA node ID (1-40) [2]:
Enter DATA node IP address by which it is known to management nodes.
Enter DATA node IP address: 10.13.20.81
```

```
-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net
Generated configuration is {0s,2sm,0m,2d} of {0s,2sm,0m,2d}
```

```
-----
Generating configuration files
```

```
Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
```

11. After reviewing the information you entered for each node, accept the configuration files.

```
View (v), accept (a), or reject (r) configuration files: a

WARNING: No such directory: /opt/JNPRshare/install/junCluster
```

12. (Optional) If prompted to create a shared directory for collecting and distributing the shared configuration of the cluster, enter **y**.

```
It is recommended that you create a shared directory owned by hadm
for the purposes of collecting and distributing shared configuration
among the various nodes of your clusters. Create this directory now,
along with the OS user account hadm if required? [y]: y
```

```
Writing shared configuration to /opt/JNPRshare/install/junCluster
```



```
-----
SBR 8.60.50006 cluster
```

```
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net
is not configured and processes are down, needs to be configured
-----
```

1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

READY: last operation succeeded, generated cluster definition.

Enter the number of the desired configuration task or quit (2,q):

13. Proceed to [“Configuring the SBRC Software on the SM Nodes” on page 463.](#)

## Configuring the SBRC Software on the SM Nodes

To configure the newly installed software on the SM nodes:

1. Configure the first SM node. When the main configuration menu is displayed, enter **3** to configure the node.

```

-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net
is not configured and processes are down, needs to be configured
-----

Configuring Cluster Node...

Enter SBR cluster name [uranus-js]: junCluster

Reading shared configuration from /opt/JNPRshare/install/junCluster

Generating configuration files

Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files: a

WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y

Cleaning directories
/opt/JNPRhadm

Applying configuration

Initializing Session State Register, please wait a few minutes...

Do you want to configure Java Runtime Environment for JDBC Feature [n] :

Please enter backup or radius directory from which to migrate.
Enter n for new configuration, s to search, or q to quit
[n]:

Enter initial admin user (UNIX account must have a valid password) [root]:
Enable Centralized Configuration Management (CCM) for this SBR node? [n]:
Configuring for use with generic database
Do you want to enable "Radius WatchDog" Process? [n]:
Do you want to enable LCI? [n]:
Do you want to configure for use with Oracle? [n]:
Removing oracle references from startup script
Do you want to configure for use with SIGTRAN? [n]:

```

```

Removing SIGTRAN references from startup script
Do you want to configure SNMP? [n]:
Configuring Admin GUI Webserver
Compatible Java version found :
Do you want to install custom SSL certificate for Admin WebServer? [n]:
Enable (e), disable (d), or preserve (p) autoboot scripts [e]:

```

```

The SBR Admin Web GUI can be launched using the following URL:
https://<servername>:2909

```

```

-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net(sm)
is configured and processes are down, may be reconfigured if desired
-----

```

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.
2. Generate Cluster Definition  
Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.
3. Configure Cluster Node  
To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.
4. Reconfigure RADIUS Server  
Only on SBR nodes, updates the existing SBR configuration.
5. Create Temporary Cluster  
Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.
6. Upgrade From Restricted Cluster License  
Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

```

READY: last operation succeeded, node configured.
Enter the number of the desired configuration task or quit (4,q): q

[07:33:44][root@uranus-js:/opt/JNPRsbr/radius/install]$

```

2. At the main menu, enter **q** to quit.

3. To configure the second SM node, log in to the node as root and execute:

**NOTE:** If you previously made changes to the **config.ini** file for your installation, before configuring the second SM and D nodes, you need to copy the **config.ini** file from first SM node's share folder path to the share folder on the other nodes (/opt/JNPRshare/install/junCluster).

root@sbrha-3: /opt/JNPRsbr/radius/install ./configure

```
Configuring SBR Software
```

```
END USER LICENSE AGREEMENT READ THIS END USER LICENSE AGREEMENT .....
.
.
.
2 Juniper Networks End User License Agreement rev. 07Aug08

Do you accept the terms in the license agreement? [n]: y
```

4. Enter **y** to accept the terms of the license agreement.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-3.englab.juniper.net
is not configured and processes are down, needs to be configured
-----

1. Unconfigure Cluster Node
   Not used when merely updating existing cluster definitions.

2. Generate Cluster Definition
   Creates new or updates existing cluster definitions.
   Modifies the shared directory but does not modify this node.

3. Configure Cluster Node
   To be preceded by 'Generate Cluster Definition' on one node.
   Must be invoked on each and every node of the cluster.
```

#### 4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

#### 5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

#### 6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q): **3**

5. When the main menu displays, enter **3** to configure the node.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-3.englab.juniper.net
is not configured and processes are down, needs to be configured
-----

Configuring Cluster Node...

Enter SBR cluster name [sbrha]: junCluster
Reading shared configuration from /opt/JNPRshare/install/junCluster

Generating configuration files

Reviewing configuration files
/opt/JNPRsbr/radius/install/tmp/config.ini
/opt/JNPRsbr/radius/install/tmp/my.cnf
/opt/JNPRsbr/radius/install/tmp/dbclusterndb.gen
View (v), accept (a), or reject (r) configuration files: a
```

6. Enter **a** to accept the configuration files.

```
WARNING: You are about to make irreversible changes to this node.
Are you sure that you wish to continue? (y,n): y
```

7. Enter **y** to make the changes to the node.
8. Press Enter through the remaining questions to preserve the node configuration.

```
Cleaning directories
/opt/JNPRhadm

Applying configuration

Initializing Session State Register, please wait a few minutes...
2447 Killed

Do you want to configure Java Runtime Environment for JDBC Feature [n] :

Please enter backup or radius directory from which to migrate.
Enter n for new configuration, s to search, or q to quit
[n]:

Enter initial admin user (UNIX account must have a valid password) [root]:
Enable Centralized Configuration Management (CCM) for this SBR node? [n]:
Configuring for use with generic database
Do you want to enable "Radius WatchDog" Process? [n]:
Do you want to enable LCI? [n]:
Do you want to configure for use with Oracle? [n]:
Removing oracle references from startup script
Do you want to configure for use with SIGTRAN? [n]:
Removing SIGTRAN references from startup script
Do you want to configure SNMP? [n]:
Configuring Admin GUI Webserver
Compatible Java version found :
Do you want to install custom SSL certificate for Admin WebServer? [n]:
Enable (e), disable (d), or preserve (p) autoboot scripts [e]:

The SBR Admin Web GUI can be launched using the following URL:
https://<servername>:2909

-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-3.englab.juniper.net(sm)
```

is configured and processes are down, may be reconfigured if desired

-----

1. Unconfigure Cluster Node  
Not used when merely updating existing cluster definitions.
2. Generate Cluster Definition  
Creates new or updates existing cluster definitions.  
Modifies the shared directory but does not modify this node.
3. Configure Cluster Node  
To be preceded by 'Generate Cluster Definition' on one node.  
Must be invoked on each and every node of the cluster.
4. Reconfigure RADIUS Server  
Only on SBR nodes, updates the existing SBR configuration.
5. Create Temporary Cluster  
Used to approximate a cluster using only this one machine.  
Intended for migration and demonstration purposes only.
6. Upgrade From Restricted Cluster License  
Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

READY: last operation succeeded, node configured.

Enter the number of the desired configuration task or quit (4,q): q

root@sbrha-3:/opt/JNPRsbr/radius/install \$

9. When the main menu displays, enter **q** to quit.

## Configuring the SBRC Software on the Data Nodes

To configure the new software on the D nodes:

1. Log in to the first D node as root and execute:

```
root@sbrha-9: /opt/JNPRsbr/radius/install$ ./configure
```

## Configuring SBR Software

```

END USER LICENSE AGREEMENT READ THIS END USER LICENSE AGREEMENT .....
.
.
.
(Translation: The parties confirm that this Agreement and all related
documentation is and will be in the English language).

2 Juniper Networks End User License Agreement rev. 07Aug08

Do you accept the terms in the license agreement? [n]: y

```

2. Enter **y** to accept the license agreement.

```

-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-9.englab.juniper.net
is not configured and processes are down, needs to be configured
-----

1.  Unconfigure Cluster Node
    Not used when merely updating existing cluster definitions.

2.  Generate Cluster Definition
    Creates new or updates existing cluster definitions.
    Modifies the shared directory but does not modify this node.

3.  Configure Cluster Node
    To be preceded by 'Generate Cluster Definition' on one node.
    Must be invoked on each and every node of the cluster.

4.  Reconfigure RADIUS Server
    Only on SBR nodes, updates the existing SBR configuration.

5.  Create Temporary Cluster
    Used to approximate a cluster using only this one machine.
    Intended for migration and demonstration purposes only.

6.  Upgrade From Restricted Cluster License

```



Used to upgrade from restricted cluster to regular cluster.  
Removes database restriction on the number of concurrent  
sessions and enables the addition of an expansion kit license

Enter the number of the desired configuration task or quit (2,q): **3**

3. When the main menu displays, enter **3** to configure the D node.

```
-----
SBR 8.60.50006 cluster
on SunOS 5.10 Generic_141444-09 node sbrha-9.englab.juniper.net
is not configured and processes are down, needs to be configured
-----
```

Configuring Cluster Node...

Enter SBR cluster name [sbrha]: **junCluster**

4. Enter the new cluster ID; in this example, "junCluster".

Reading shared configuration from /opt/JNPRshare/install/junCluster

Generating configuration files

Reviewing configuration files

/opt/JNPRsbr/radius/install/tmp/config.ini

/opt/JNPRsbr/radius/install/tmp/my.cnf

View (v), accept (a), or reject (r) configuration files: **a**

5. Enter **a** to accept the configuration files.

WARNING: You are about to make irreversible changes to this node.

Are you sure that you wish to continue? (y,n): **y**

6. Enter **y** to make the changes to the node.

Cleaning directories

/opt/JNPRhadm

Applying configuration

Initializing Session State Register, please wait a few minutes...

```
-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-9.englab.juniper.net(d)
is configured and processes are down, may be reconfigured if desired
-----
```

1. Unconfigure Cluster Node

Not used when merely updating existing cluster definitions.

2. Generate Cluster Definition

Creates new or updates existing cluster definitions.

Modifies the shared directory but does not modify this node.

3. Configure Cluster Node

To be preceded by 'Generate Cluster Definition' on one node.

Must be invoked on each and every node of the cluster.

4. Reconfigure RADIUS Server

Only on SBR nodes, updates the existing SBR configuration.

5. Create Temporary Cluster

Used to approximate a cluster using only this one machine.

Intended for migration and demonstration purposes only.

6. Upgrade From Restricted Cluster License

Used to upgrade from restricted cluster to regular cluster.

Removes database restriction on the number of concurrent sessions and enables the addition of an expansion kit license

READY: last operation succeeded, node configured.

Enter the number of the desired configuration task or quit (2,q): **q**

root@sbrha-9:/opt/JNPRsbr/radius/install\$

7. When the main menu displays, enter **q** to quit.
8. Repeat Step 1 through Step 7 on the next data node.

## Starting the SSR Processes

Start the SSR process on each node in the cluster, one at a time. Start the SSR process on the nodes in the following order:

- SM nodes
- D nodes

After you start the SSR process on the node, check the status of the node and wait for the process to finish starting before moving on to the next node.

1. Log in to the first SM node.

Execute:

```
root@uranus-js: /opt/JNPRsbr/radius/sbrd start ssr
```

```
Starting ssr management processes
MySQL Cluster Management Server mysql-5.7.25 ndb-7.6.9
-Node 51: connected (Version 7.1.8)

Starting ssr auxiliary processes
root@uranus-js:
```

2. After starting the SSR process on the node, monitor the status of the node and make sure that it reconnects to the cluster.

Execute:

```
root@uranus-js: /opt/JNPRsbr/radius ./sbrd status
```

```
-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net(sm)
-----

Connected to Management Server at: 10.13.20.89:5235
```

```

[ndbd(NDB)]      2 node(s)
id=1 (not connected, accepting connect from 10.13.20.83)
id=2 (not connected, accepting connect from 10.13.20.81)

[ndb_mgmd(MGM)]  2 node(s)
id=51 @10.13.20.89 (mysql-5.6.36 ndb-7.4.15)
id=52 (not connected, accepting connect from 10.13.20.77)

[mysqld(API)]    4 node(s)
id=61 (not connected, accepting connect from 10.13.20.89)
id=62 (not connected, accepting connect from 10.13.20.77)
id=100 (not connected, accepting connect from 10.13.20.89)
id=101 (not connected, accepting connect from 10.13.20.77)

      hadm  8216 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
      hadm  8195 ndb_mgmd --config-cache=0
      hadm  8281 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
--datadir=/o

[07:47:02][root@uranus-js:/opt/JNPRsbr/radius]$

```

### 3. On the second SM:

Execute:

```
root@sbrha-3: /opt/JNPRsbr/radius/sbrd start ssr
```

```
root@sbrha-3: /opt/JNPRsbr/radius/sbrd status
```

### 4. Start the SSR process on the first D node.

```

20:57:47][root@sbrha-9:/opt/JNPRsbr/radius]$ ./sbrd start ssr
Starting ssr data processes
You have mail.
2010-10-21 20:57:50 [ndbd] INFO      -- Angel connected to '10.13.20.89:5235'
2010-10-21 20:57:50 [ndbd] INFO      -- Angel allocated nodeid: 1
Initializing ssr
[20:57:50][root@sbrha-9:/opt/JNPRsbr/radius

```

### 5. After starting the SSR process on the node, monitor the status of the node and make sure that it reconnects to the cluster.

Execute:

```
root@sbrha-9: /opt/JNPRsbr/radius ./sbrd status
```

```
-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node sbrha-9.englab.juniper.net(d)
-----
```

```
root 6025 ndbmt d --initial
--ndb-connectstring=nodeid=1;10.13.20.89:5235,10.13.20.77:5235
root 6026 ndbmt d --initial
--ndb-connectstring=nodeid=1;10.13.20.89:5235,10.13.20.77:5235
```

6. On the second D node:

Execute:

```
root@sbrha-7: /opt/JNPRsbr/radius/sbrd start ssr
```

```
root@sbrha-7: /opt/JNPRsbr/radius/sbrd status
```

7. After the SSR process has been started on all four nodes, check the status of the cluster from one of the SM nodes.

Execute:

```
root@uranus-js: /opt/JNPRsbr/radius ./sbrd status
```

The status of the cluster should look similar to the following example. Notice that all nodes are connected except for (in this example) node IDs 100 and 101. These will be connected after you re-create and restore the cluster database and start the RADIUS processes.

```
[07:49:09][root@uranus-js:/opt/JNPRsbr/radius]$ ./sbrd status
```

```
-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net(sm)
-----
```

```
Connected to Management Server at: 10.13.20.89:5235
```

```
[ndbd(NDB)]      2 node(s)
id=1      @10.13.20.83  (mysql-5.6.36 ndb-7.4.15, Nodegroup: 0, Master)
id=2      @10.13.20.81  (mysql-5.6.36 ndb-7.4.15, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=51     @10.13.20.89  (mysql-5.6.36 ndb-7.4.15)
id=52     @10.13.20.77  (mysql-5.6.36 ndb-7.4.15)
```

```
[mysqld(API)]    4 node(s)
id=61   @10.13.20.89   (mysql-5.6.36 ndb-7.4.15)
id=62   @10.13.20.77   (mysql-5.6.36 ndb-7.4.15)
id=100  (not connected, accepting connect from 10.13.20.89)
id=101  (not connected, accepting connect from 10.13.20.77)

      hadm  8216 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
      hadm  8195 ndb_mgmd --config-cache=0
      hadm  8281 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
--datadir=/o

[07:49:48][root@uranus-js:/opt/JNPRsbr/radius]$
```

## Creating and Restoring the SSR Database

Create and restore the SSR database.

**NOTE:** If you need to customize the sessions database to match your existing cluster session database, see [“Customizing the SSR Database Current Sessions Table” on page 144](#). Any customization must be done prior to running the **CreateDB.sh** script.

1. Create the database by logging in to one of the SM nodes as hadm.

Execute:

**hadm@uranus-js:~>./CreateDB.sh**

```
Creating database "SteelBeltedRadius" (using ENGINE ndbcluster).
Creating misc tables.
Creating IP Pool, Range and Address tables.
Creating Wimax tables.
Creating SIM Auth tables.
Creating TTLS Resumption Cache tables.
Creating Current Sessions table.
Proxy AutoStop feature configured.
Session Timeout on Missed Account Stop feature configured.
Use Single Class Attribute feature configured.
```

```
Creating User Concurrency table.
Creating stored routines.
```

2. Restore the database from the backup you made in [“Backing Up the State of the Sessions in Your SSR Database”](#) on page 450.

Execute:

```
hadm@uranus-js:~>mysql -D SteelBeltedRadius </opt/backup_sbr.sql
```

3. Execute the **UpdateSchema.pl** utility for schema upgrades.
  - a. First, determine whether your schema needs to be updated for this release by executing **UpdateSchema** with no arguments as the hadm user. This will return a list of schemas which need to be updated for the installed version of software.

Execute:

```
hadm@uranus-js:~>perl ./UpdateSchema.pl
```

```
Usage: perl UpdateSchema.pl [version]
       where [version] is 5.4, 5.5, 7.2.5, or 8.4 and above
       Other versions do not need a schema update.
```

From SBR 8.5 release, to increase the Column size of AcctSessionId from 24 to 48, the below usage is recommended:

```
Usage: perl UpdateSchema.pl <To_Version> ColumnUpdate:Sbr_AcctSessionId
```

**NOTE:** Future SBR Carrier software releases may require schema updates, we recommend you execute this step every time you upgrade to determine if a schema upgrade is required.

- b. If you are running a software release that requires a schema change, execute the command again with the version to which you are upgrading. For example, if you are upgrading to SBR Carrier Release 8.6.0:

Execute:

```
hadm@uranus-js:~>perl ./UpdateSchema.pl 8.6 ColumnUpdate:Sbr_AcctSessionId
```

```
Altering Tables for adding Sbr_Ipv6Address- Please wait...Done
Altering Tables for changing Sbr_AcctSessionId size from 24 to 48- Please
wait... Done
```

**NOTE:** The time required for altering the CST depends on the number of sessions stored in the CST. For example, SBR Carrier takes approximately 30 minutes to update a CST that contains 5 million sessions.

4. Check the database restore.

Execute:

```
hadm@uranus-js:~>./ShowSessions.sh -c
```

```
Total Sessions: 611443
```

5. If you need to customize the sessions database, see [“Customizing the SSR Database Current Sessions Table” on page 144.](#)
6. Configure at least one IP address pool and one range using the SSR Administration Scripts. See [“Testing the Installation with DemoSetup.sh” on page 482.](#) See also *Session State Register Administration* in the *SBR Carrier Administration and Configuration Guide*.

## Starting the RADIUS Process

Start the RADIUS process on each S and SM node.

To start the RADIUS process on each SM node:

1. Log in to the SM node as root.

Execute:

```
root@uranus-js: ./sbrd start radius
```

```
Starting radius server processes
RADIUS: Process ID of daemon is 8993
RADIUS: Starting DCF system
RADIUS: Configuration checksum: 59 3C 6B AA
radius started
[08:06:38][root@uranus-js:/opt/JNPRsbr/radius]$
```

2. Repeat Step 1 on the second SM node.
3. Check the status of the cluster from each SM node.

Execute:



**root@uranus-js: ./sbrd status**

The cluster status should look similar to the following example. Notice that all nodes are now connected to the cluster.

```
-----
SBR 8.60.50006 cluster junCluster{0s,2sm,0m,2d}
on SunOS 5.10 Generic_141444-09 node uranus-js.englab.juniper.net(sm)
-----

Connected to Management Server at: 10.13.20.89:5235

[ndbd(NDB)]      2 node(s)
id=1      @10.13.20.83  (mysql-5.6.36 ndb-7.4.15, Nodegroup: 0, Master)
id=2      @10.13.20.81  (mysql-5.6.36 ndb-7.4.15, Nodegroup: 0)

[ndb_mgmd(MGM)] 2 node(s)
id=51     @10.13.20.89  (mysql-5.6.36 ndb-7.4.15)
id=52     @10.13.20.77  (mysql-5.6.36 ndb-7.4.15)

[mysqld(API)]   4 node(s)
id=61     @10.13.20.89  (mysql-5.6.36 ndb-7.4.15)
id=62     @10.13.20.77  (mysql-5.6.36 ndb-7.4.15)
id=100    @10.13.20.89  (mysql-5.6.36 ndb-7.4.15)
id=101    @10.13.20.77  (mysql-5.6.36 ndb-7.4.15)

10.13.20.89.1646          Idle
10.13.20.89.1813          Idle
10.13.20.89.1645          Idle
10.13.20.89.1812          Idle
      *.1812              *.*              0      0 49152      0 LISTEN
      *.1813              *.*              0      0 49152      0 LISTEN

      hadm  8216 /bin/sh /opt/JNPRmysql/install/bin/mysqld_safe
      hadm  8195 ndb_mgmd --config-cache=0
      hadm  8281 /opt/JNPRmysql/install/bin/mysqld --basedir=/opt/JNPRmysql/install
--datadir=/o
      root  8993 radius sbr.xml

[08:09:44][root@uranus-js:/opt/JNPRsbr/radius]$
```

## Launching Web GUI

Now that the RADIUS process is running, you can complete the configuration using Web GUI. For details, see [“Basic SBR Carrier Node Configuration” on page 138](#). For complete details, see the *SBR Carrier Administration and Configuration Guide*.

**PART 9**

# Appendix

---

This part contains the following appendix:

## APPENDIX A

## Testing the Installation with DemoSetup.sh

---

**IN THIS SECTION**

- [Creating a Test Database | 482](#)
- [DemoSetup.sh Syntax | 484](#)

This appendix describes how to set up a test database to test the installation. Creation of the test database requires at least one IP address pool to be configured for SBR Carrier. These topics are in this appendix:

## Creating a Test Database

If all the nodes come up correctly and see each other, you can create a test database with the **DemoSetup.sh** script.

To create a test database, on each management node:

1. Log in as **root**. (You use both root and hadm accounts; using two windows makes this easier.)
2. Stop the SBR Carrier node.
  - a. Change directories to **/opt/JNPRsbr/radius/**.
  - b. Execute:  
**sbrd stop radius**
3. Log in (using the second window) as **hadm**.

4. Create a database with the **DemoSetup.sh** script, using the default values.

Execute: **DemoSetup.sh**

For example:

**DemoSetup.sh**

5. Respond to the script prompts with **yes**.

For example:

```
SBRs must be offline, do you want to proceed? <yes|no> yes
This will destroy the " SteelBeltedRadius" database (if it exists), OK? <yes|no>
yes
```

6. In the hadm window, check the status of all nodes as they come online.

**/opt/JNPRsbr/radius/sbrd status**


Results similar to this example are displayed:

```
hadmUser$>/opt/JNPRsbr/radius/sbrd status
```

```
[ndbd(NDB)]      2 node(s)
id=10   @172.28.84.163  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0, Master)
id=11   @172.28.84.113  (mysql-5.7.25 ndb-7.6.9, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=1     @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=2     @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
```

```
[mysqld(API)]    4 node(s)
id=21   @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=22   @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
id=30   @172.28.84.36  (mysql-5.7.25 ndb-7.6.9)
id=31   @172.28.84.166 (mysql-5.7.25 ndb-7.6.9)
```



Examine the line starting with “id=”, and verify that there are no references to **starting**, **connecting**, or **not connected**. Any of these references indicate the process has either not finished starting, or the node is not connected properly. You may need to execute the **sbrd status** command more than once because it only shows a snapshot of activity; the display does not refresh. Do not proceed to the next node until you are sure the process has started properly and the node is connected.

# DemoSetup.sh Syntax

The **DemoSetup.sh** script creates a new database with the name specified in the **DBName.txt** file, creates all required tables, and populates the database with a sample IP address pool configuration. You can use this program to test the installation and to become familiar with Session State Register.

## Syntax

DemoSetup.sh [ *numpools maxranges minaddrs maxaddrs* ]  
 DemoSetup.sh -h

## Options

Table 48: DemoSetup.sh Options

Option	Description
<i>numpools</i>	<p>A number in the range 1–25 that specifies the number of IP address pools to create.</p> <p>Default value is 5.</p>
<i>maxranges</i>	<p>A number in the range 1–20 that specifies the number of address ranges within each IP address pool to create.</p> <p>Default value is 4.</p>

Table 48: DemoSetup.sh Options (continued)

Option	Description
<i>minaddrs</i>	A number in the range 1–100,000 that specifies the minimum number of IP addresses per range to create.  Default value is 1000.
<i>maxaddrs</i>	A number in the range 1–100,000 that specifies the maximum number of IP addresses per range to create.  Default value is the value of <i>minaddrs</i> x 2.
-h	Displays help for the DemoSetup script.

## Notes

The script must be executed by `hadm` on a management node.

## Example

This example displays the contents of the **DBName.txt** file (to verify the name (**SteelBeltedRadius**) to be used to create the database) and creates a database called **SteelBeltedRadius**.

```
hadm$- DemoSetup.sh
SBRs must be offline, do you want to proceed? <yes|no> yes
This will destroy the " SteelBeltedRadius" database (if it exists),
OK? <yes|no> yes
      NUMPOOLS=5 MAXRANGES=4 MINADDRS=1000 MAXADDRS=2000
      Destroying old database
      Database "SteelBeltedRadius" destroyed.
      Creating new database
      Creating database "SteelBeltedRadius" (using ENGINE ndbcluster).
      Creating misc tables.
      Creating IP Pool, Range and Address tables.
      Creating Current Sessions table.
          (Proxy AutoStop feature not configured.)
          (Session Timeout on Missed Account Stop feature not configured.)
```

```

        (Use Single Class Attribute feature not configured.)
Creating User Concurrency table.
Creating stored routines.
Adding pool: A-PLATINUM
Adding range: A-PLATINUM 238.203.131.14 670
Adding range: A-PLATINUM 48.226.119.162 761
Adding pool: B-GOLD
Adding range: B-GOLD 90.169.221.242 549
Adding range: B-GOLD 20.97.26.189 89
Adding range: B-GOLD 114.23.180.47 407
Adding range: B-GOLD 236.99.7.33 384
Adding pool: C-SILVER
Adding range: C-SILVER 201.213.13.198 978
Adding range: C-SILVER 15.28.1.17 22
Adding range: C-SILVER 28.189.195.246 2
Adding range: C-SILVER 132.185.74.85 11
Adding pool: D-BRONZE
Adding range: D-BRONZE 122.219.182.131 1247
Adding range: D-BRONZE 247.168.228.227 20
Adding pool: E-ZINC
Adding range: E-ZINC 135.54.24.139 1794
IpPools:
+-----+-----+-----+
| Name| Ord| Count|
+-----+-----+-----+
| (z o m b i e)| 0| 0|
| A-PLATINUM| 1| 1,431|
| B-GOLD| 2| 1,429|
| C-SILVER| 3| 1,013|
| D-BRONZE| 4| 1,267|
| E-ZINC| 5| 1,794|
+-----+-----+-----+
Total Pools : 5 + 1 zombie pseudo-pool.
Total Ranges: 13 + 0 zombie pseudo-ranges.
Total Addrs : 6,934 + 0 zombie pseudo-addrs.
IpRanges:
+-----+-----+-----+-----+
| Pool| StartAddr| EndAddr| Count|
+-----+-----+-----+-----+
| (z o m b i e) | (v a r i o u s)| (v a r i o u s)| 0|
| A-PLATINUM| 48.226.119.162| 48.226.122.154| 761|
| A-PLATINUM| 238.203.131. 14| 238.203.133.171| 670|
| B-GOLD| 20. 97. 26.189| 20. 97. 27. 21| 89|
| B-GOLD| 90.169.221.242| 90.169.224. 22| 549|

```



```
| B-GOLD| 114. 23.180. 47| 114. 23.181.197| 407|  
| B-GOLD| 236. 99. 7. 33| 236. 99. 8.160| 384|  
| C-SILVER| 15. 28. 0. 17| 15. 28. 0. 38| 22|  
| C-SILVER| 28.189.195.246| 28.189.195.247| 2|  
| C-SILVER| 132.185. 74. 85| 132.185. 74. 95| 11|  
| C-SILVER| 201.213. 13.198| 201.213. 17.151| 978|  
| D-BRONZE| 122.219.182.131| 122.219.187. 97| 1,247|  
| D-BRONZE| 247.168.228.227| 247.168.228.246| 20|  
| E-ZINC| 135. 54. 24.139| 135. 54. 31.140| 1,794|  
+-----+-----+-----+-----+
```