

PROGRESS[®] ARTIX[™]

WSDL Extension Reference

Version 5.6, December 2011

© 2011 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved.

These materials and all Progress[®] software products are copyrighted and all rights are reserved by Progress Software Corporation. The information in these materials is subject to change without notice, and Progress Software Corporation assumes no responsibility for any errors that may appear therein. The references in these materials to specific platforms supported are subject to change.

Actional, Apama, Artix, Business Empowerment, DataDirect (and design), DataDirect Connect, DataDirect Connect64, DataDirect Technologies, DataDirect XML Converters, DataDirect XQuery, DataXtend, Dynamic Routing Architecture, EdgeXtend, Empowerment Center, Fathom, Fuse Mediation Router, Fuse Message Broker, Fuse Services Framework, IntelliStream, IONA, Making Software Work Together, Mindreef, ObjectStore, OpenEdge, Orbix, PeerDirect, POSSENET, Powered by Progress, PowerTier, Progress, Progress DataXtend, Progress Dynamics, Progress Business Empowerment, Progress Empowerment Center, Progress Empowerment Program, Progress OpenEdge, Progress Profiles, Progress Results, Progress Software Developers Network, Progress Sonic, ProVision, PS Select, Savvion, SequeLink, Shadow, SOAPscope, SOAPStation, Sonic, Sonic ESB, SonicMQ, Sonic Orchestration Server, SpeedScript, Stylus Studio, Technical Empowerment, Web-Speed, Xcalia (and design), and Your Software, Our Technology—Experience the Connection are registered trademarks of Progress Software Corporation or one of its affiliates or subsidiaries in the U.S. and/or other countries. AccelEvent, Apama Dashboard Studio, Apama Event Manager, Apama Event Modeler, Apama Event Store, Apama Risk Firewall, AppsAlive, AppServer, ASPen, ASP-in-a-Box, BusinessEdge, Business Making Progress, Cache-Forward, CloudEdge, DataDirect Spy, DataDirect SupportLink, Fuse, FuseSource, Future Proof, GVAC, High Performance Integration, ObjectStore Inspector, ObjectStore Performance Expert, OpenAccess, Orbacus, Pantero, POSSE, ProDataSet, Progress Arcade, Progress CloudEdge, Progress Control Tower, Progress ESP Event Manager, Progress ESP Event Modeler, Progress Event Engine, Progress RFID, Progress RPM, Progress Software Business Making Progress, PSE Pro, SectorAlliance, SeeThinkAct, Shadow z/Services, Shadow z/Direct, Shadow z/Events, Shadow z/Presentation, Shadow Studio, SmartBrowser, SmartComponent, SmartDataBrowser, SmartDataObjects, SmartDataView, SmartDialog, SmartFolder, SmartFrame, SmartObjects, SmartPanel, SmartQuery, SmartViewer, SmartWindow, Sonic Business Integration Suite, Sonic Process Manager, Sonic Collaboration Server, Sonic Continuous Availability Architecture, Sonic Database Service, Sonic Workbench, Sonic XML Server, The Brains Behind BAM, WebClient, and Who Makes Progress are trademarks or service marks of Progress Software Corporation and/or its subsidiaries or affiliates in the U.S. and other countries. Java is a registered trademark of Oracle and/or its affiliates. Any other marks contained herein may be trademarks of their respective owners.

Third Party Acknowledgments:

Progress Artix ESB for C++ v5.6 incorporates Xalan v2.3.1 technologies from the Apache Software Foundation (<http://www.apache.org>). Such Apache technologies are subject to the following terms and conditions: The Apache Software License, Version 1.1. Copyright (C) 1999-2002 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 "Ant", "Xerces," "Xalan," "Log 4J," 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. 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. For more information on the Apache Software Foundation, please see <http://www.apache.org/>. Xalan was originally based on software copyright (c) 1999, Lotus Development Corporation., <http://www.lotus.com>. Xerces was originally based on software copyright (c) 1999, International Business Machines, Inc., <http://www.ibm.com>.

Progress Artix ESB for C++ v5.6 incorporates Xerces C++ v2.4 technology from the Apache Software Foundation (<http://www.apache.org>). Such Apache technology is subject to the following terms and conditions: The Apache Software License, Version 1.1 - Copyright (c) 1999-2001 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.

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.

Progress Artix ESB for C++ v5.6 incorporates Apache Xerces v2.5.0 technology from the Apache Software Foundation (http://www.apache.org). Such Apache technology is subject to the following terms and conditions: The Apache Software License, Version 1.1 - Copyright (c) 1999-2002 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.

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 (c) 1999, International Business Machines, Inc., <http://www.ibm.com>. For more information on the Apache Software Foundation, please see <http://www.apache.org/>.

Progress Artix ESB for C++ v5.6 incorporates Xerces C++ v1.7 technology from the Apache Software Foundation (<http://www.apache.org>). Such Apache technology is subject to the following terms and conditions: The Apache Software License, Version 1.1. - Copyright (c) 1999-2004 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 "Xalan" 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.
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

ING 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 (c) 1999, Lotus Development Corporation., <http://www.lotus.com>. For more information on the Apache Software Foundation, please see <http://www.apache.org/>.

Progress Artix ESB for C++ v5.6 incorporates Apache Velocity v1.3 technology from the Apache Software Foundation (<http://www.apache.org>). Such Apache technology is subject to the following terms and conditions: The Apache Software License, Version 1.1 - Copyright (c) 2000-2003 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 acknowledgement: "This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>)." Alternately, this acknowledgement may appear in the software itself, if and wherever such third-party acknowledgements normally appear.
4. The names "The Jakarta Project", "Velocity", 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.
5. Products derived from this software may not be called "Apache", "Velocity" nor may "Apache" appear in their names without prior written permission of the Apache Group.

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.

Progress Artix ESB for C++ v5.6 incorporates Log4J v1.2.6 technology from the Apache Software Foundation (<http://www.apache.org>). Such Apache technology is subject to the following terms and conditions: The Apache Software License, Version 1.1 - Copyright (C) 1999 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 "log4j" 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.

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. For more information on the Apache Software Foundation, please see <http://www.apache.org/>.

(a) Progress Artix ESB for C++ v5.6 incorporates JDOM Beta 9 technology from JDOM. Such technology is subject to the following terms and conditions: Copyright (C) 2000-2004 Jason Hunter & Brett McLaughlin. 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 disclaimer that follows these conditions in the documentation and/or other materials provided with the distribution. 3. The name "JDOM" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact request_AT_jdom_DOT_org. 4. Products derived from this software may not be called "JDOM", nor may "JDOM" appear in their name, without prior written permission from the JDOM Project Management request_AT_jdom_DOT_org. In addition, we request (but do not require) that you include in the end-user documentation provided with the redistribution and/or in the software itself an acknowledgement equivalent to the following: "This

product includes software developed by the JDOM Project (<http://www.jdom.org/>).¹ Alternatively, the acknowledgment may be graphical using the logos available at <http://www.jdom.org/images/logos>. 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 JDOM AUTHORS OR THE PROJECT 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 JDOM Project and was originally created by Jason Hunter <jhunter_AT_jdom_DOT_org> and Brett McLaughlin <brett_AT_jdom_DOT_org>. For more information on the JDOM Project, please see <<http://www.jdom.org/>>

Progress Artix ESB for C++ v5.6 incorporates IBM-ICU v2.6 and IBM-ICU v2.6.1 technologies from IBM. Such technologies are subject to the following terms and conditions: Copyright (c) 1995-2003 International Business Machines Corporation and others All rights reserved. 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, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation. 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 OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM 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. Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder. All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Progress Artix ESB for C++ v5.6 incorporates John Wilson MinML v1.7 technology from John Wilson. Such technology is subject to the following terms and conditions: Copyright (c) 1999, John Wilson (tug@wilson.co.uk). All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright

notice, this list of conditions and the following disclaimer. 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. All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by John Wilson. The name of John Wilson may not be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY JOHN WILSON ``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 JOHN WILSON 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.

Progress Artix ESB for C++ v5.6 incorporates SourceForge - NET-SNMP v5.0.7 technology from SourceForge and Networks Associates Technology, Inc. Such technology is subject to the following terms and conditions: Various copyrights apply to this package, listed in various separate parts below. Please make sure that you read all the parts. Up until 2001, the project was based at UC Davis, and the first part covers all code written during this time. From 2001 onwards, the project has been based at SourceForge, and Networks Associates Technology, Inc hold the copyright on behalf of the wider Net-SNMP community, covering all derivative work done since then. An additional copyright section has been added as Part 3 below also under a BSD license for the work contributed by Cambridge Broadband Ltd. to the project since 2001. An additional copyright section has been added as Part 4 below also under a BSD license for the work contributed by Sun Microsystems, Inc. to the project since 2003. Code has been contributed to this project by many people over the years it has been in development, and a full list of contributors can be found in the README file under the THANKS section. ---- Part 1: CMU/UCD copyright notice: (BSD like) ---- Copyright 1989, 1991, 1992 by Carnegie Mellon University. Derivative Work - 1996, 1998-2000. Copyright 1996, 1998-2000 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. ----
Part 2: Networks Associates Technology, Inc copyright notice (BSD) ----- Copyright (c)
2001-2003, 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: *Redistributions of source code must retain the above
copyright notice, this list of conditions and the following disclaimer.* Redistributions in
binary form must reproduce the above copyright notice, this list of conditions and the fol-
lowing disclaimer in the documentation and/or other materials provided with the distribu-
tion.* 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 COPY-
RIGHT 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 PUR-
POSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR
CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPE-
CIAL, 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 LIA-
BILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY
WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSI-
BILITY OF SUCH DAMAGE. ---- Part 3: Cambridge Broadband Ltd. copyright notice
(BSD) ----- Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd.
All rights reserved. Redistribution and use in source and binary forms, with or without mod-
ification, are permitted provided that the following conditions are met:*Redistributions of
source code must retain the above copyright notice, this list of conditions and the following
disclaimer.* 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.* The name of Cambridge Broadband Ltd. may not be used to
endorse or promote products derived from this software without specific prior written per-
mission. 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 COPY-
RIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPE-
CIAL, 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 LIA-
BILITY, 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. ---- Part 4: Sun
Microsystems, Inc. copyright notice (BSD) -----Copyright © 2003 Sun Microsystems, Inc.,
4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved. Use is
subject to license terms below. This distribution may include materials developed by third
parties. Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trade-
marks of Sun Microsystems, Inc. in the U.S. and other countries. Redistribution and use in
source and binary forms, with or without modification, are permitted provided that the fol-

lowing conditions are met:* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.* 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.* Neither the name of the Sun Microsystems, 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. ---- Part 5: Sparta, Inc copyright notice (BSD) -----Copyright (c) 2003-2005, Sparta, 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:* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.* 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.* Neither the name of Sparta, 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. ---- Part 6: Cisco/BUPTNIC copyright notice (BSD) ----- Copyright (c) 2004, Cisco, Inc and Information Network Center of Beijing University of Posts and Telecommunications. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. * 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. * Neither the name of Cisco, Inc, Beijing University of Posts and Telecommunications, nor the names of their 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. ---- Part 7: Fabasoft R&D Software GmbH & Co KG copyright notice (BSD) ----- Copyright (c) Fabasoft R&D Software GmbH & Co KG, 2003 oss@fabasoft.com Author: Bernhard Penz. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.* 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. * The name of Fabasoft R&D Software GmbH & Co KG or any of its subsidiaries, brand or product names 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.

Progress Artix ESB for C++ v5.6 incorporates OpenSSL/SSLey v0.9.8i technology from OpenSSL.org. Such Technology is subject to the following terms and conditions: LICENSE ISSUES =====

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLey license apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.

OpenSSL License -----

/*

=====
=====

Copyright (c) 1998-2008 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.
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 (ey@cryptsoft.com).
This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLey License -----

Copyright (C) 1995-1998 Eric Young (ey@cryptsoft.com) All rights reserved.

This package is an SSL implementation written by Eric Young (ey@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 adhered 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 routines 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.]

Progress Artix ESB for C++ v5.6 incorporates Bouncycastle v1.3.3 cryptographic technology from the Legion Of The Bouncy Castle (<http://www.bouncycastle.org>). Such Bouncycastle 1.3.3 cryptographic technology is subject to the following terms and conditions: Copyright (c) 2000 - 2006 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.

Progress Artix ESB for C++ v5.6 incorporates PCRE 7.8 from PCRE for the purpose of providing a set of functions that implement regular expression pattern matching using the same syntax and semantics as Perl 5. Such technology is subject to the following terms and conditions: PCRE LICENCE. PCRE is a library of functions to support regular expressions whose syntax and semantics are as close as possible to those of the Perl 5 language. Release 7 of PCRE is distributed under the terms of the "BSD" licence, as specified below. The documentation for PCRE, supplied in the "doc" directory, is distributed under the same terms as the software itself. The basic library functions are written in C and are freestanding. Also included in the distribution is a set of C++ wrapper functions. THE BASIC LIBRARY FUNCTIONS. Written by: Philip Hazel. Email local part: ph10. Email domain: cam.ac.uk. University of Cambridge Computing Service, Cambridge, England. Copyright (c) 1997-2008 University of Cambridge All rights reserved. THE C++ WRAPPER FUNCTIONS. Contributed by: Google Inc. Copyright (c) 2007-2008, Google Inc. All rights reserved. THE "BSD" LICENCE. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. * 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. * Neither the name of the University of Cambridge nor the name of Google Inc. nor the names of their 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 OWNER 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.

Progress Artix ESB for C++ v5.6 incorporates mcpp v2.6.4 from Kiyoshi Matsui. Such technology is subject to the following terms and conditions: Copyright (c) 1998, 2002-2007 Kiyoshi Matsui kmatsui@t3.rim.or.jp All rights reserved. This software including the files in this directory is provided under the following license. 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 ``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 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.

Progress Artix ESB for C++ v5.6 contains IBM Licensed Materials Copyright IBM Corporation 2010 (IBM 32-bit Runtime Environment for AIX, Java Technology Edition v 1.6.0 SR9 FP2).

Updated: December 5, 2011

Contents

List of Tables	25
List of Figures	27
Preface	29
What is Covered in this Book	29
Who Should Read this Book	29
How to Use this Book	29
The Artix Documentation Library	29

Part I Bindings

SOAP 1.1 Binding	33
Runtime Compatibility	33
soap:binding	33
soap:operation	34
soap:body	35
soap:header	37
soap:fault	38
SOAP 1.2 Binding	41
Runtime Compatibility	41
wsoap12:binding	41
wsoap12:operation	42
wsoap12:body	43
wsoap12:header	45
wsoap12:fault	46
MIME Multipart/Related Binding	49
Runtime Compatibility	49
Namespace	49

mime:multipartRelated	50
mime:part	50
mime:content	50
CORBA Binding and Type Map	53
CORBA Binding Extension Elements	54
Runtime Compatibility	54
C++ Runtime Namespace	54
Primitive Type Mapping	54
corba:binding	56
corba:operation	56
corba:param	57
corba:return	58
corba:raises	58
.....	59
corba:typeMapping	59
corba:struct	60
corba:member	60
corba:enum	61
corba:enumerator	62
corba:fixed	62
corba:union	64
corba:unionbranch	64
corba:case	65
corba:alias	66
corba:array	67
corba:sequence	68
corba:exception	69
corba:anonsequence	70
corba:anonstring	72
corba:object	73
Tuxedo FML Binding	79
Runtime Compatibility	79
Namespace	79
FML\XMLSchema Support	80
tuxedo:binding	80
tuxedo:fieldTable	80

tuxedo:field	81
tuxedo:operation	81
Fixed Binding	83
Runtime Compatibility	83
Namespace	83
fixed:binding	83
fixed:operation	84
fixed:body	84
fixed:field	85
fixed:enumeration	88
fixed:choice	90
fixed:case	90
fixed:sequence	92
Tagged Binding	95
Runtime Compatibility	95
Namespace	95
tagged:binding	95
tagged:operation	97
tagged:body	97
tagged:field	98
tagged:enumeration	98
tagged:sequence	99
tagged:choice	101
tagged:case	102
TibrvMsg Binding	105
Runtime Compatibility	105
Namespace	105
TIBRVMSG to XMLSchema Type Mapping	106
tibrv:binding	107
tibrv:operation	108
tibrv:input	109
tibrv:output	110
tibrv:array	111
tibrv:msg	115
tibrv:field	116

tibrv:context	117
Chapter 9 XML Binding	119
Runtime Compatibility	119
Namespace	119
xformat:binding	120
xformat:body	120
Pass Through Binding	121
Runtime Compatibility	121
Namespace	121
tagged:binding	122
Part II Ports	
HTTP Port	125
Standard WSDL Elements	126
http:address	126
soap:address	126
wsoap12:address	126
Configuration Extensions for C++	127
Namespace	127
http-conf:client	127
http-conf:server	130
Attribute Details	133
AuthorizationType	133
Authorization	133
Accept	133
AcceptLanguage	134
AcceptEncoding	135
ContentType	135
ContentEncoding	136
Host	136
Connection	137
CacheControl	137
BrowserType	140

Referer	140
ProxyServer	142
ProxyAuthorizationType	142
ProxyAuthorization	142
UseSecureSockets	143
RedirectURL	143
ServerCertificateChain	143
CORBA Port	145
Runtime Compatibility	145
C++ Runtime Namespace	145
corba:address	145
corba:policy	146
IIOP Tunnel Port	149
Runtime Compatibility	149
Namespace	149
iiop:address	149
iiop:payload	150
iiop:policy	151
Chapter 15 WebSphere MQ Port	153
Artix Extension Elements	154
Runtime Compatibility	154
Namespace	154
mq:client	154
mq:server	157
Attribute Details	160
Server_Client	160
AliasQueueName	161
UsageStyle	163
CorrelationStyle	164
AccessMode	165
MessagePriority	166
Delivery	166
Transactional	167
ReportOption	168
Format	170

Tuxedo Port	173
Runtime Compatibility	173
Namespace	173
tuxedo:server	173
tuxedo:service	174
tuxedo:input	174
JMS Port	175
C++ Runtime Extensions	176
Namespace	176
jms:address	176
jms:JMSNamingProperty	177
jms:client	178
jms:server	178
Tibco/Rendezvous Port	181
Artix Extension Elements	182
Runtime Compatibility	182
Namespace	182
tibrv:port	183
Attribute Details	187
bindingType	187
callbackLevel	187
responseDispatchTimeout	188
transportService	188
transportNetwork	188
cmTransportServerName	188
cmQueueTransportServerName	189
File Transfer Protocol Port	191
Runtime Compatibility	191
Namespace	191
ftp:port	192
ftp:properties	192
ftp:property	193

Part III Other Extensions

Routing	197
Runtime Compatibility	197
Namespace	197
routing:expression	197
routing:route	198
routing:source	198
routing:query	199
routing:destination	199
routing:transportAttribute	200
routing:equals	201
routing:greater	202
routing:less	202
routing:startswith	203
routing:endswith	203
routing:contains	204
routing:empty	204
routing:nonempty	205
Transport Attribute Context Names	205
Security	207
Runtime Compatibility	207
Namespace	207
bus-security:security	207
Codeset Conversion	211
Runtime Compatibility	211
Namespace	211
i18n-context:client	211
i18n-context:server	212
Index	213

CONTENTS

List of Tables

Table 1: Primitive Type Mapping for CORBA Plug-in	54
Table 2: Complex IDL Type Mappings	59
Table 3: FML Type Support	80
Table 4: Values of tuxedo:fieldTable Element's type Attribute	81
Table 5: TIBCO to XMLSchema Type Mapping	106
Table 6: Functions Used for Specifying TibrvMsg Array Element Names	113
Table 7: Settings for http-conf:client CacheControl	137
Table 8: Settings for http-conf:server CacheControl	139
Table 9: Server_Client Attribute Settings	160
Table 10: <i>UsageStyle Settings</i>	163
Table 11: <i>MQGET and MQPUT Actions</i>	164
Table 12: <i>Artix WebSphere MQ Access Modes</i>	165
Table 13: <i>TabDelivery Attribute Settings</i>	167
Table 14: <i>Transactional Attribute Settings</i>	167
Table 15: <i>ReportOption Attribute Settings</i>	168
Table 16: <i>FormatType Attribute Settings</i>	170
Table 17: TIB/RV Supported Payload formats	187
Table 18: Context QNames	205

LIST OF TABLES

List of Figures

Figure 1: *MQ Remote Queues*

162

LIST OF FIGURES

Preface

What is Covered in this Book

This book is a reference to all of the Artix ESB specific WSDL extensions used in Artix contracts.

Who Should Read this Book

This book is intended for Artix users who are familiar with Artix concepts including:

- WSDL
- XMLSchema
- Artix interface design

In addition, this book assumes that the reader is familiar with the transports and middleware implementations with which they are working.

How to Use this Book

This book contains the following parts:

- [“Bindings”](#)—contains descriptions for all the WSDL extensions used to define the payload formats supported by Artix.
- [“Ports”](#)—contains descriptions for all the WSDL extensions used to define the transports supported by Artix.
- [“Other Extensions”](#)—contains descriptions for the WSDL extensions used by Artix to support features like routing.

The Artix Documentation Library

For information on the organization of the Artix library, the document conventions used, and where to find additional resources, see [Using the Artix Library](#).

Part I

Bindings

In this part

This part contains the following chapters:

SOAP 1.1 Binding	page 33
SOAP 1.2 Binding	page 41
MIME Multipart/Related Binding	page 49
CORBA Binding and Type Map	page 53
Tuxedo FML Binding	page 79
Fixed Binding	page 83
Tagged Binding	page 95
TibrvMsg Binding	page 105
XML Binding	page 119
Pass Through Binding	page 121

SOAP 1.1 Binding

This chapter describes the extensions used to define a SOAP 1.1 message.

Runtime Compatibility

The SOAP binding is defined by a standard set of WSDL extensors.

soap:binding

Synopsis

```
<soap:binding style="..." transport="..." />
```

Description

The `soap:binding` element specifies that the payload format to use is a SOAP 1.1 message. It is a child of the WSDL `binding` element.

Attributes

The following attributes are defined within the `soap:binding` element.

- [style](#)
- [transport](#)

style

The value of the `style` attribute within the `soap:binding` element acts as the default for the `style` attribute within each `soap:operation` element. It indicates whether request/response operations within this binding are RPC-based (that is, messages contain parameters and return values) or document-based (that is, messages contain one or more documents).

Valid values are `rpc` and `document`. The specified value determines how the SOAP `Body` element within a SOAP message is structured.

If `rpc` is specified, each message part within the SOAP `Body` element is a parameter or return value and will appear inside a wrapper element within the SOAP `Body` element. The name of the wrapper element must match the operation name. The namespace of the wrapper element is based on the value of the `soap:body namespace` attribute. The message parts within the wrapper element correspond to operation parameters and must appear in the same order as the parameters in the operation. Each part name must match the parameter name to which it corresponds.

For example, the SOAP `Body` element of a SOAP request message is as follows if the style is RPC-based:

```
<SOAP-ENV:Body>
  <m:GetStudentGrade xmlns:m="URL">
    <StudentCode>815637</StudentCode>
    <Subject>History</Subject>
  </m:GetStudentGrade>
</SOAP-ENV:Envelope>
```

If `document` is specified, message parts within the SOAP `Body` element appear directly under the SOAP `Body` element as body entries and do not appear inside a wrapper element that corresponds to an operation. For example, the SOAP `Body` element of a SOAP request message is as follows if the style is document-based:

```
<SOAP-ENV:Body>
  <StudentCode>815637</StudentCode>
  <Subject>History</Subject>
</SOAP-ENV:Envelope>
```

transport

The `transport` attribute defaults to the URL that corresponds to the HTTP binding in the W3C SOAP specification (<http://schemas.xmlsoap.org/soap/http>). If you want to use another transport (for example, SMTP), modify this value as appropriate for the transport you want to use.

soap:operation

Synopsis

```
<soap:operation style="..." soapAction="..." />
```

Description

The `soap:operation` element is a child of the WSDL `operation` element. A `soap:operation` element is used to encompass information for an operation as a whole, in terms of input criteria, output criteria, and fault information.

Attributes

The following attributes are defined within a `soap:operation` element:

- [style](#)
- [soapAction](#)

style

This indicates whether the relevant operation is RPC-based (that is, messages contain parameters and return values) or document-based (that is, messages contain one or more documents).

Valid values are `rpc` and `document`. The default value for `soap:operation style` is based on the value specified for the `soap:binding style` attribute.

See “[style](#)” on page 33 for more details of the `style` attribute.

soapAction

This specifies the value of the `SOAPAction` HTTP header field for the relevant operation. The value must take the form of the absolute URI that is to be used to specify the intent of the SOAP message.

Note: This attribute is mandatory only if you want to use SOAP over HTTP. Leave it blank if you want to use SOAP over any other transport.

soap:body

Synopsis

```
<soap:body use="..." encodingStyle="..." namespace="..."
parts="..." />
```

Description

The `soap:body` element in a binding is a child of the `input`, `output`, and `fault` child elements of the WSDL `operation` element. A `soap:body` element is used to provide information on how message parts are to appear inside the body of a SOAP message. As explained in “[soap:operation](#)” on page 34, the structure of the SOAP `Body` element within a SOAP message is dependent on the setting of the `soap:operation style` attribute.

Attributes

The following attributes are defined within a `soap:body` element:

- [use](#)
- [encodingStyle](#)
- [namespace](#)
- [parts](#)

use

This mandatory attribute indicates how message parts are used to denote data types. Each message part relates to a particular data type that in turn might relate to an abstract type definition or a concrete schema definition.

An abstract type definition is a type that is defined in some remote encoding schema whose location is referenced in the WSDL contract via an `encodingStyle` attribute. In this case, types are serialized based on the set of rules defined by the specified encoding style.

A concrete schema definition relates to types that are defined in the WSDL contract itself, within a `schema` element within the `types` component of the contract.

The following are valid values for the `use` attribute:

- `encoded`
- `literal`

If `encoded` is specified, the `type` attribute that is specified for each message part (within the `message` component of the WSDL contract) is used to reference an abstract type defined in some remote encoding schema. In this case, a concrete SOAP message is produced by applying encoding rules to the abstract types. The encoding rules are based on the encoding style identified in the `soap:body encodingStyle` attribute. The encoding takes as input the `name` and `type` attribute for each message part (defined in the `message` component of the WSDL contract). If the encoding style allows variation in the message format for a given set of abstract types, the receiver of the message must ensure they can understand all the format variations.

If `literal` is specified, either the `element` or `type` attribute that is specified for each message part (within the `message` component of the WSDL contract) is used to reference a concrete schema definition (defined within the `types` component of the WSDL contract). If the `element` attribute is used to reference a concrete schema definition, the referenced element in the SOAP message appears directly under the `SOAP Body` element (if the operation style is document-based) or under a part accessor element that has the same name as the message part (if the operation style is RPC-based). If the `type` attribute is used to reference a concrete schema definition, the referenced type in the SOAP message becomes the schema type of the `SOAP Body` element (if the operation style is document-based) or of the part accessor element (if the operation style is document-based).

encodingStyle

This attribute is used when the `soap:body use` attribute is set to `encoded`. It specifies a list of URIs (each separated by a space) that represent encoding styles that are to be used within the SOAP message. The URIs should be listed in order, from the most restrictive encoding to the least restrictive.

This attribute can also be used when the `soap:body use` attribute is set to `literal`, to indicate that a particular encoding was used to derive the concrete format, but that only the specified variation is supported. In this case, the sender of the SOAP message must conform exactly to the specified schema.

namespace

If the `soap:operation style` attribute is set to `rpc`, each message part within the SOAP `Body` element of a SOAP message is a parameter or return value and will appear inside a wrapper element within the SOAP `Body` element. The name of the wrapper element must match the operation name. The namespace of the wrapper element is based on the value of the `soap:body namespace` attribute.

parts

This attribute is a space separated list of parts from the parent `input`, `output`, or `fault` element. When `parts` is set, only the specified parts of the message are included in the SOAP `Body` element. The unlisted parts are not transmitted unless they are placed into the SOAP header.

soap:header

Synopsis

```
<soap:header message="..." part="..." use="..." encodingStyle="..." namespace="..." />
```

Description

The `soap:header` element in a binding is an optional child of the `input`, `output`, and `fault` elements of the WSDL `operation` element. A `soap:header` element defines the information that is placed in a SOAP header element. You can define any number of `soap:header` elements for an operation. As explained in [“soap:operation” on page 34](#), the structure of the SOAP header within a SOAP message is dependent on the setting of the `soap:operation` element’s `style` attribute.

Attributes

The `soap:header` element has the following attributes.

<code>message</code>	Specifies the qualified name of the message from which the contents of the SOAP header is taken.
<code>part</code>	Specifies the name of the message part that is placed into the SOAP header.
<code>use</code>	Used in the same way as the <code>use</code> attribute within the <code>soap:body</code> element. See “ use ” on page 36 for more details.
<code>encodingStyle</code>	Used in the same way as the <code>encodingStyle</code> attribute within the <code>soap:body</code> element. See “ encodingStyle ” on page 37 for more details.
<code>namespace</code>	If the <code>soap:operation style</code> attribute is set to <code>rpc</code> , each message part within the SOAP header of a SOAP message is a parameter or return value and will appear inside a wrapper element within the SOAP header. The name of the wrapper element must match the operation name. The namespace of the wrapper element is based on the value of the <code>soap:header namespace</code> attribute.

soap:fault**Synopsis**

```
<soap:fault name="..." use="..." encodingStyle="..." />
```

Description

The `soap:fault` element is a child of the WSDL `fault` element within an operation component. Only one `soap:fault` element is defined for a particular operation. The operation must be a request-response or solicit-response type of operation, with both `input` and `output` elements. The `soap:fault` element is used to transmit error and status information within a SOAP response message.

Note: A fault message must consist of only a single message part. Also, it is assumed that the `soap:operation` element’s `style` attribute is set to `document`, because faults do not contain parameters.

Attributes

The `soap:fault` element has the following attributes:

<code>name</code>	Specifies the name of the fault. This relates back to the <code>name</code> attribute for the <code>fault</code> element specified for the corresponding operation within the <code>portType</code> component of the WSDL contract.
<code>use</code>	This attribute is used in the same way as the <code>use</code> attribute within the <code>soap:body</code> element. See “use” on page 36 for more details.
<code>encodingStyle</code>	This attribute is used in the same way as the <code>encodingStyle</code> attribute within the <code>soap:body</code> element. See “encodingStyle” on page 37 for more details.

SOAP 1.2 Binding

This chapter describes the extensions used to define a SOAP 1.2 message.

Runtime Compatibility

The SOAP 1.2 binding is defined by a standard set of WSDL extensors.

wsoap12:binding

Synopsis

```
<wsoap12:binding style="..." transport="..." />
```

Description

The `wsoap12:binding` element specifies that the payload format to use is a SOAP 1.2 message. It is a child of the WSDL `binding` element.

Attributes

The following attributes are defined within the `wsoap12:binding` element.

- [style](#)
- [transport](#)

style

The value of the `style` attribute acts as the default for the `style` attribute within each `wsoap12:operation` element. It indicates whether request/response operations within this binding are RPC-based (that is, messages contain parameters and return values) or document-based (that is, messages contain one or more documents).

Valid values are `rpc` and `document`. The specified value determines how the SOAP `Body` element within a SOAP message is structured.

If `rpc` is specified, each message part within the SOAP `Body` element is a parameter or return value and will appear inside a wrapper element within the SOAP `Body` element. The name of the wrapper element must match the operation name. The namespace of the wrapper element is based on the value of the `soap:body namespace` attribute. The message parts within the wrapper element correspond to operation parameters and must appear in the same order as the parameters in the operation. Each part name must match the parameter name to which it corresponds.

For example, the SOAP `Body` element of a SOAP request message is as follows if the style is RPC-based:

```
<SOAP-ENV:Body>
  <m:GetStudentGrade xmlns:m="URL">
    <StudentCode>815637</StudentCode>
    <Subject>History</Subject>
  </m:GetStudentGrade>
</SOAP-ENV:Envelope>
```

If `document` is specified, message parts within the SOAP `Body` element appear directly under the SOAP `Body` element as body entries and do not appear inside a wrapper element that corresponds to an operation. For example, the SOAP `Body` element of a SOAP request message is as follows if the style is document-based:

```
<SOAP-ENV:Body>
  <StudentCode>815637</StudentCode>
  <Subject>History</Subject>
</SOAP-ENV:Envelope>
```

transport

The `transport` attribute specifies a URL describing the SOAP transport to which this binding corresponds. The URL that corresponds to the HTTP binding in the W3C SOAP specification is `http://schemas.xmlsoap.org/soap/http`. If you want to use another transport (for example, SMTP), modify this value as appropriate for the transport you want to use.

wsoap12:operation

Synopsis

```
<wsoap12:operation style="..." soapAction="..."
soapActionRequired="..." />
```

Description

The `wsoap12:operation` element is a child of the WSDL `operation` element. A `soap:operation` element is used to encompass information for an operation as a whole, in terms of input criteria, output criteria, and fault information.

Attributes

The following attributes are defined within a `wsoap12:operation` element:

- [style](#)
- [soapAction](#)
- [soapActionRequired](#)

style

This indicates whether the relevant operation is RPC-based (that is, messages contain parameters and return values) or document-based (that is, messages contain one or more documents).

Valid values are `rpc` and `document`. The default value for the `wsoap12:operation` element's `style` attribute is based on the value specified for the [wsoap12:binding](#) element's `style` attribute.

soapAction

This specifies the value of the `SOAPAction` HTTP header field for the relevant operation. The value must take the form of the absolute URI that is to be used to specify the intent of the SOAP message.

Note: This attribute is mandatory only if you want to use SOAP 1.2 over HTTP. Leave it blank if you want to use SOAP 1.2 over any other transport.

soapActionRequired

The `soapActionRequired` is a boolean that specifies if the value of the [soapAction](#) attribute must be conveyed in the request message. When the value of `soapActionRequired` is `true`, the [soapAction](#) attribute must be present. The default is to `true`.

wsoap12:body

Synopsis

```
<wsoap12:body use="..." encodingStyle="..." namespace="..."  
parts="..." />
```

Description

The `wsoap12:body` element in a binding is a child of the input, output, and fault child elements of the WSDL `operation` element. A `wsoap12:body` element is used to provide information on how message parts are to appear inside the body of a SOAP 1.2 message. As explained in [“wsoap12:operation” on page 42](#), the

structure of the SOAP `Body` element within a SOAP message is dependent on the setting of the `soap:operation style` attribute.

Attributes

The following attributes are defined within a `wsoap12:body` element:

- [use](#)
- [encodingStyle](#)
- [namespace](#)
- [parts](#)

use

This mandatory attribute indicates how message parts are used to denote data types. Each message part relates to a particular data type that in turn might relate to an abstract type definition or a concrete schema definition.

An abstract type definition is a type that is defined in some remote encoding schema whose location is referenced in the WSDL contract via an `encodingStyle` attribute. In this case, types are serialized based on the set of rules defined by the specified encoding style.

A concrete schema definition relates to types that are defined in the WSDL contract itself, within a `schema` element within the `types` component of the contract.

The following are valid values for the `use` attribute:

- `literal`
- `encoded`

Note: Artix 4.1 does not support encoded messages when using SOAP 1.2.

If `literal` is specified, either the `element` or `type` attribute that is specified for each message part (within the `message` component of the WSDL contract) is used to reference a concrete schema definition (defined within the `types` component of the WSDL contract). If the `element` attribute is used to reference a concrete schema definition, the referenced element in the SOAP 1.2 message appears directly under the SOAP `Body` element (if the operation style is document-based) or under a part accessor element that has the same name as the message part (if the operation style is RPC-based). If the `type` attribute is used to reference a concrete schema definition, the referenced type in the SOAP 1.2 message becomes the schema type of the SOAP `Body` element (if the operation style is document-based) or of the part accessor element (if the operation style is document-based).

encodingStyle

This attribute is only used when the `wsoap12:body` element's `use` attribute is set to `encoded`, and the `wsoap12:binding` element's `style` attribute is set to `rpc`. It specifies the URI that represents the encoding rules that used to construct the SOAP 1.2 message.

namespace

If the `soap:operation` element's `style` attribute is set to `rpc`, each message part within the SOAP `Body` element of a SOAP 1.2 message is a parameter or return value and will appear inside a wrapper element within the SOAP `Body` element. The name of the wrapper element must match the operation name. The namespace of the wrapper element is based on the value of the `soap:body namespace` attribute.

parts

This attribute is a space separated list of parts from the parent `input`, `output`, or `fault` element. When the `parts` attribute is set, only the specified parts of the message are included in the SOAP `Body` element. The unlisted parts are not transmitted unless they are placed into the SOAP header.

wsoap12:header

Synopsis

```
<wsoap12:header message="..." part="..." use="..."  
encodingStyle="..." namespace="..." />
```

Description

The `wsoap12:header` element in a binding is an optional child of the `input`, `output`, and `fault` elements of the WSDL `operation` element. A `wsoap12:header` element defines the information that is placed in a SOAP 1.2 header element. You can define any number of `wsoap12:header` elements for an operation. As explained in [“wsoap12:operation” on page 42](#), the structure of the header within a SOAP 1.2 message is dependent on the setting of the `wsoap12:operation` element's `style` attribute.

Attributes

The `wsoap12:header` element has the following attributes.

<code>message</code>	Specifies the qualified name of the message from which the contents of the SOAP header is taken.
<code>part</code>	Specifies the name of the message part that is placed into the SOAP header.

<code>use</code>	Used in the same way as the <code>wsoap12:body</code> element's <code>use</code> attribute.
<code>encodingStyle</code>	Used in the same way as the <code>wsoap12:body</code> element's <code>encodingStyle</code> attribute.
<code>namespace</code>	Specifies the namespace to be assigned to the header element when the <code>use</code> attribute is set to <code>encoded</code> . The header is constructed in all cases as if the <code>wsoap12:binding</code> element's <code>style</code> attribute had a value of <code>document</code> .

wsoap12:fault

Synopsis

```
<wsoap12:fault name="..." namespace="..." use="..."
encodingStyle="..." />
```

Description

The `wsoap12:fault` element is a child of the WSDL `fault` element within a WSDL `operation` element. The operation must have both `input` and `output` elements. The `wsoap12:fault` element is used to transmit error details and status information within a SOAP 1.2 response message.

Note: A fault message must consist of only a single message part. Also, it is assumed that the `wsoap12:operation` element's `style` attribute is set to `document`, because faults do not contain parameters.

Attributes

The `wsoap12:fault` element has the following attributes:

<code>name</code>	Specifies the name of the fault. This relates back to the <code>name</code> attribute for the <code>fault</code> element specified for the corresponding operation within the <code>portType</code> component of the WSDL contract.
<code>namespace</code>	Specifies the namespace to be assigned to the wrapper element for the fault. This attribute is ignored if the <code>style</code> attribute of either the <code>wsoap12:binding</code> element of the containing binding or of the <code>wsoap12:operation</code> element of the containing operation is either omitted or has a value of <code>document</code> . This attribute is required if the value of the <code>wsoap12:binding</code> element's <code>style</code> attribute is set to <code>rpc</code> .
<code>use</code>	This attribute is used in the same way as the <code>wsoap12:body</code> element's <code>use</code> attribute.
<code>encodingStyle</code>	This attribute is used in the same way as the <code>wsoap12:body</code> element's <code>encodingStyle</code> attribute

MIME

Multipart/Related Binding

This chapter describes the extensions that are used to define a SOAP message binding that contains binary data.

Runtime Compatibility

The MIME extensions are defined by a standard.

Namespace

The WSDL extensions used to define the MIME multipart/related messages are defined in the namespace `http://schemas.xmlsoap.org/wsdl/mime/`.

In the discussion that follows, it is assumed that this namespace is prefixed with `mime`. The entry in the WSDL `defintion` element to set this up is shown in [Example 1](#).

Example 1: *MIME Namespace Specification in a Contract*

```
xmlns:mime="http://schemas.xmlsoap.org/wsdl/mime/"
```

mime:multipartRelated

Synopsis

```
<mime:multipartRelated>
  <mime:part ...>
  ...
</mime:part>
...
</mime:multipartRelated>
```

Description

The `mime:multipartRelated` element is the child of an `input` element or an `output` element that is part of a SOAP binding. It tells Artix that the message body is going to be a multipart message that potentially contains binary data.

`mime:multipartRelated` elements in Artix contain one or more [mime:part](#) elements that describe the individual parts of the message.

mime:part

Synopsis

```
<mime:part name="...">
  ...
</mime:part>
```

Description

The `mime:part` element is the child of a [mime:multipartRelated](#) element. It is used to define the parts of a multi-part message. The first `mime:part` element must contain the [soap:body](#) element or the [wsoap12:body](#) element that would normally appear in a SOAP binding. The remaining `mime:part` elements define the attachments that are being sent in the message using a [mime:content](#) element.

Attributes

The `mime:part` element has a single attribute called `name`. `name` is a unique string that is used to identify the part being described.

mime:content

Synopsis

```
<mime:content part="..." type="..." />
```

Description

The `mime:content` element is the child of a [mime:part](#) element. It defines the binary content being passed as an attachment to a SOAP message.

Attributes

The `mime:content` element has the following attributes:

`part` Specifies the name of the WSDL `part` element, from the parent message definition, that is used as the content of this part of the MIME multipart message being placed on the wire.

`type` Specifies the MIME type of the data in this message part. MIME types are defined as a type and a subtype using the syntax *type/subtype*.

There are a number of predefined MIME types such as `image/jpeg` and `text/plain`. The MIME types are maintained by IANA and described in the following:

- *Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies*
(<ftp://ftp.isi.edu/in-notes/rfc2045.txt>)
- *Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types*
(<ftp://ftp.isi.edu/in-notes/rfc2046.txt>).

CORBA Binding and Type Map

Artix CORBA support uses a combination of a WSDL binding element and a corba:typeMapping element to unambiguously define CORBA Messages.

In this chapter

This chapter discusses the following topics:

CORBA Binding Extension Elements	page 54
Type Map Extension Elements	page 59

CORBA Binding Extension Elements

Runtime Compatibility

The CORBA binding extensions are compatible with the C++ runtime.

C++ Runtime Namespace

The WSDL extensions used for the C++ Runtime CORBA binding and the CORBA data mappings are defined in the namespace

`http://schemas.iona.com/bindings/corba`. The Artix designer adds the following namespace declaration to any contract that uses the C++ runtime CORBA binding:

```
xmlns:corba="http://schemas.iona.com/bindings/corba"
```

Primitive Type Mapping

Most primitive IDL types are directly mapped to primitive XML Schema types. [Table 1](#) lists the mappings for the supported IDL primitive types.

Table 1: *Primitive Type Mapping for CORBA Plug-in*

IDL Type	XML Schema Type	CORBA Binding Type	Artix C++ Type
Any	xsd:anyType	corba:any	IT_Bus::AnyHolder
boolean	xsd:boolean	corba:boolean	IT_Bus::Boolean
char	xsd:byte	corba:char	IT_Bus::Char
wchar	xsd:string	corba:wchar	
double	xsd:double	corba:double	IT_Bus::Double
float	xsd:float	corba:float	IT_Bus::Float
octet	xsd:unsignedByte	corba:octet	IT_Bus::Octet

Table 1: *Primitive Type Mapping for CORBA Plug-in*

IDL Type	XML Schema Type	CORBA Binding Type	Artix C++ Type
long	xsd:int	corba:long	IT_Bus::Long
long long	xsd:long	corba:longlong	IT_Bus::LongLong
short	xsd:short	corba:short	IT_Bus::Short
string	xsd:string	corba:string	IT_Bus::String
wstring	xsd:string	corba:wstring	
unsigned short	xsd:unsignedShort	corba:ushort	IT_Bus::UShort
unsigned long	xsd:unsignedInt	corba:ulong	IT_Bus::ULong
unsigned long long	xsd:unsignedLong	corba:ulonglong	IT_Bus::ULongLong
Object	wsa:EndpointReferenceType	corba:object	WS_Addresssing::EndpointReferenceType
TimeBase::UtcT	xsd:dateTime ^a	corba:dateTime	IT_Bus::DateTime

a. The mapping between `xsd:dateTime` and `TimeBase::UtcT` is only partial. For the restrictions see [“Unsupported time/date values” on page 55](#)

Unsupported types

The following CORBA types are not supported:

- long double
- Value types
- Boxed values
- Local interfaces
- Abstract interfaces
- Forward-declared interfaces

Unsupported time/date values

The following `xsd:dateTime` values cannot be mapped to `TimeBase::UtcT`:

- Values with a local time zone. Local time is treated as a 0 UTC time zone offset.
- Values prior to 15 October 1582.
- Values greater than approximately 30,000 A.D.

The following `TimeBase::UtcT` values cannot be mapped to `xsd:dateTime`:

- Values with a non-zero `inaccco` or `inacchi`.
- Values with a time zone offset that is not divisible by 30 minutes.
- Values with time zone offsets greater than 14:30 or less than -14:30.
- Values with greater than millisecond accuracy.
- Values with years greater than 9999.

corba:binding

Synopsis

```
<corba:binding repositoryID="..." bases=".." />
```

Description

The `corba:binding` element indicates that the binding is a CORBA binding.

Attributes

This element has two attributes:

<code>repositoryID</code>	A required attribute whose value is the full type ID of the CORBA interface. The type ID is embedded in an object's IOR and must conform to the format <code>IDL:module/interface:1.0.</code>
<code>bases</code>	An optional attribute whose value is the type ID of the interface from which the interface being bound inherits.

Examples

For example, the following IDL:

```
//IDL
interface clash{};
interface bad : clash{};
```

would produce the following `corba:binding`:

```
<corba:binding repositoryID="IDL:bad:1.0"
bases="IDL:clash:1.0"/>
```

corba:operation

Synopsis

```
<corba:operation name="..." >
  <corba:param ... />
  ...
  <corba:return ... />
  <corba:raises ... />
```

```
</corba:operation>
```

Description

The `corba:operation` element is a child element of the WSDL `operation` element and describes the parts of the operation's messages. It has one or more of the following children:

- [corba:param](#)
- [corba:return](#)
- [corba:raises](#)

Attributes

The `corba:operation` attribute takes a single attribute, `name`, which duplicates the name given in `operation`.

corba:param**Synopsis**

```
<corba:param name="..." mode="..." idltype="..." />
```

Description

The `corba:param` element is a child of `corba:operation`. Each `part` element of the input and output messages specified in the logical operation, except for the part representing the return value of the operation, must have a corresponding `corba:param` element. The parameter order defined in the binding must match the order specified in the IDL definition of the operation.

Attributes

The `corba:param` element has the following required attributes:

<code>mode</code>	Specifies the direction of the parameter. The values directly correspond to the IDL directions: <code>in</code> , <code>inout</code> , <code>out</code> . Parameters set to <code>in</code> must be included in the input message of the logical operation. Parameters set to <code>out</code> must be included in the output message of the logical operation. Parameters set to <code>inout</code> must appear in both the input and output messages of the logical operation.
<code>idltype</code>	Specifies the IDL type of the parameter. The type names are prefaced with <code>corba:</code> for primitive IDL types, and <code>corbatm:</code> for complex data types, which are mapped out in the <code>corba:typeMapping</code> portion of the contract. See “Type Map Extension Elements” on page 59 .
<code>name</code>	Specifies the name of the parameter as given in the <code>name</code> attribute of the corresponding <code>part</code> element.

corba:return

Synopsis

```
<corba:return name="..." idltype="..." />
```

Description

The `corba:return` element is a child of `corba:operation` and specifies the return type, if any, of the operation.

Attributes

The `corba:return` element has two attributes:

<code>name</code>	Specifies the name of the parameter as given in the logical portion of the contract.
<code>idltype</code>	Specifies the IDL type of the parameter. The type names are prefaced with <code>corba:</code> for primitive IDL types and <code>corbatm:</code> for complex data types which are mapped out in the <code>corba:typeMapping</code> portion of the contract.

corba:raises

Synopsis

```
<corba:raises exception="..." />
```

Description

The `corba:raises` element is a child of `corba:operation` and describes any exceptions the operation can raise. The exceptions are defined as fault messages in the logical definition of the operation. Each fault message must have a corresponding `corba:raises` element.

Attributes

The `corba:raises` element has one required attribute, `exception`, which specifies the type of data returned in the exception.

Type Map Extension Elements

corba:typeMapping

Synopsis

```
<corba:typeMapping
targetNamespace="http://schemas.ionac.com/bindings/corba/typemap">
...
</corba:typeMapping>
```

Description

Because complex types (such as structures, arrays, and exceptions) require a more involved mapping to resolve type ambiguity, the full mapping for a complex type is described in a `corba:typeMapping` element in an Artix contract. This element contains a type map describing the metadata required to fully describe a complex type as a CORBA data type. This metadata may include the members of a structure, the bounds of an array, or the legal values of an enumeration.

Attributes

The `corba:typeMapping` element requires a `targetNamespace` attribute that specifies the namespace for the elements defined by the type map.

Examples

[Table 2](#) shows the mappings from complex IDL types to Artix CORBA types.

Table 2: *Complex IDL Type Mappings*

IDL Type	CORBA Binding Type
struct	<code>corba:struct</code>
enum	<code>corba:enum</code>
fixed	<code>corba:fixed</code>
union	<code>corba:union</code>
typedef	<code>corba:alias</code>
array	<code>corba:array</code>
sequence	<code>corba:sequence</code>
exception	<code>corba:exception</code>

corba:struct

Synopsis

```
<corba:struct name="..." type="..." repositoryID="..." />
  <corba:member ... />
  ...
</corba:struct>
```

The `corba:struct` element is used to represent XMLSchema types that are defined using `complexType` elements. The elements of the structure are described by a series of `corba:member` elements.

Attributes

A `corba:struct` element requires three attributes:

<code>name</code>	A unique identifier used to reference the CORBA type in the binding.
<code>type</code>	The logical type the structure is mapping.
<code>repositoryID</code>	The fully specified repository ID for the CORBA type.

corba:member

Synopsis

```
<corba:member name="..." idlType="..." />
```

Description

The `corba:member` element is used to define the parts of the structure represented by the parent element. The elements must be declared in the same order used in the IDL representation of the CORBA type.

Attributes

A `corba:member` requires two attributes:

<code>name</code>	The name of the element
<code>idltype</code>	The IDL type of the element. This type can be either a primitive type or another complex type that is defined in the type map.

Examples

For example, you may have a structure, `personalInfo`, similar to the one in [Example 2](#).

Example 2: *personalInfo*

```
enum hairColorType {red, brunette, blonde};

struct personalInfo
{
    string name;
    int age;
    hairColorType hairColor;
}
```

It can be represented in the CORBA type map as shown in [Example 3](#).

Example 3: *CORBA Type Map for personalInfo*

```
<corba:typeMapping targetNamespace="http://schemas.ionas.com/bindings/corba/typemap">
...
  <corba:struct name="personalInfo" type="xsd:personalInfo" repositoryID="IDL:personalInfo:1.0">
    <corba:member name="name" idltype="corba:string"/>
    <corba:member name="age" idltype="corba:long"/>
    <corba:member name="hairColor" idltype="corbatm:hairColorType"/>
  </corba:struct>
</corba:typeMapping>
```

The `idltype corbatm:hairColorType` refers to a complex type that is defined earlier in the CORBA type map.

corba:enum**Synopsis**

```
<corba:enum name="..." type="..." repositoryID="...">
  <corba:enumerator ... />
  ...
</corba:enum>
```

The `corba:enum` element is used to represent enumerations. The values for the enumeration are described by a series of [corba:enumerator](#) elements.

Attributes

A `corba:enum` element requires three attributes:

<code>name</code>	A unique identifier used to reference the CORBA type in the binding.
<code>type</code>	The logical type the structure is mapping.
<code>repositoryID</code>	The fully specified repository ID for the CORBA type.

corba:enumerator**Synopsis**

```
<corba:enumerator value="..." />
```

Description

The `corba:enumerator` element represents the values of an enumeration. The values must be listed in the same order used in the IDL that defines the CORBA enumeration.

Attributes

A `corba:enumerator` element takes one attribute, `value`.

Examples

For example, the enumeration defined in [Example 2 on page 61](#), `hairColorType`, can be represented in the CORBA type map as shown in [Example 4](#):

Example 4: *CORBA Type Map for hairColorType*

```
<corba:typeMapping targetNamespace="http://schemas.ionac.com/bindings/corba/typemap">
...
  <corba:enum name="hairColorType" type="xsd:hairColorType"
    repositoryID="IDL:hairColorType:1.0">
    <corba:enumerator value="red"/>
    <corba:enumerator value="brunette"/>
    <corba:enumerator value="blonde"/>
  </corba:enum>
</corba:typeMapping>
```

corba:fixed**Synopsis**

```
<corba:fixed name="..." repositoryID="..." type="..." digits="..."
scale="..." />
```

Description

Fixed point data types are a special case in the Artix contract mapping. A CORBA fixed type is represented in the logical portion of the contract as the XML Schema primitive type `xsd:decimal`. However, because a CORBA fixed type requires additional information to be fully mapped to a physical CORBA data type, it must

also be described in the CORBA type map section of an Artix contract using a `corba:fixed` element.

Attributes

A `corba:fixed` element requires five attributes:

<code>name</code>	A unique identifier used to reference the CORBA type in the binding.
<code>repositoryID</code>	The fully specified repository ID for the CORBA type.
<code>type</code>	The logical type the structure is mapping (for CORBA fixed types, this is always <code>xsd:decimal</code>).
<code>digits</code>	The upper limit for the total number of digits allowed. This corresponds to the first number in the fixed type definition.
<code>scale</code>	The number of digits allowed after the decimal point. This corresponds to the second number in the fixed type definition.

Examples

For example, the fixed type defined in [Example 5](#), `myFixed`, would be described

Example 5: *myFixed Fixed Type*

```
\\IDL
typedef fixed<4,2> myFixed;
```

by a type entry in the logical type description of the contract, as shown in [Example 6](#).

Example 6: *Logical description from myFixed*

```
<xsd:element name="myFixed" type="xsd:decimal"/>
```

In the CORBA type map portion of the contract, it would be described by an entry similar to [Example 7](#). Notice that the description in the CORBA type map includes the information needed to fully represent the characteristics of this particular fixed data type.

Example 7: *CORBA Type Map for myFixed*

```
<corba:typeMapping targetNamespace="http://schemas.ionac.com/bindings/corba/typemap">
...
  <corba:fixed name="myFixed" repositoryID="IDL:myFixed:1.0" type="xsd:decimal" digits="4"
    scale="2"/>
</corba:typeMapping>
```

corba:union

Synopsis

```
<corba:union name="..." type="..." discriminator="..."
    repositoryID="...">
  <corba:unionbranch ... />
  ...
</corba:union>
```

Description

The `corba:union` element is used to resolve the relationship between a union's discriminator and its members. A `corba:union` element is required for every CORBA union defined in an IDL contract. The members of the union are described using a series of nested `corba:unionbranch` elements.

Attributes

A `corba:union` element has four mandatory attributes:

<code>name</code>	A unique identifier used to reference the CORBA type in the binding.
<code>type</code>	The logical type the structure is mapping.
<code>discriminator</code>	The IDL type used as the discriminator for the union.
<code>repositoryID</code>	The fully specified repository ID for the CORBA type.

corba:unionbranch

Synopsis

```
<corba:unionbranch name="..." idltype="..." default="...">
  <corba:case ... />
  ...
</corba:unionbranch>
```

Description

The `corba:unionbranch` element defines the members of a union. Each `corba:unionbranch` except for one describing the union's default member will have at least one `corba:case` element as a child.

Attributes

A `corba:unionbranch` element has two required attributes and one optional attribute.

<code>name</code>	A unique identifier used to reference the union member.
<code>idltype</code>	The IDL type of the union member. This type can be either a primitive type or another complex type that is defined in the type map.

`default` The optional attribute specifying if this member is the default case for the union. To specify that the value is the default set this attribute to `true`.

corba:case

Synopsis

```
<corba:case label="..." />
```

Description

The `corba:case` element defines the explicit relationship between the discriminator's value and the associated union member.

Attributes

The `corba:case` element's only attribute, `label`, specifies the value used to select the union member described by the [corba:unionbranch](#).

Examples

For example consider the union, `myUnion`, shown in [Example 8](#):

Example 8: *myUnion IDL*

```
//IDL
union myUnion switch (short)
{
  case 0:
    string case0;
  case 1:
  case 2:
    float case12;
  default:
    long caseDef;
};
```

For example `myUnion`, [Example 8](#), would be described with a CORBA type map entry similar to that shown in [Example 9](#).

Example 9: *myUnion CORBA type map*

```
<corba:typeMapping
  targetNamespace="http://schemas.iona.com/bindings/corba/typemap">
  ...
  <corba:union name="myUnion" type="xsd1:myUnion"
    discriminator="corba:short" repositoryID="IDL:myUnion:1.0">
    <corba:unionbranch name="case0" idltype="corba:string">
      <corba:case label="0"/>
    </corba:unionbranch>
```

Example 9: *myUnion CORBA type map*

```
<corba:unionbranch name="case12" idltype="corba:float">
  <corba:case label="1"/>
  <corba:case label="2"/>
</corba:unionbranch>
<corba:unionbranch name="caseDef" idltype="corba:long"
  default="true"/>
</corba:union>
</corba:typeMapping>
```

corba:alias**Synopsis**

```
<corba:alias name="..." type="..." repositoryID="..." />
```

Description

The `corba:alias` element is used to represent a `typedef` statement in an IDL contract.

Attributes

The `corba:alias` element has three attributes:

<code>name</code>	The value of the <code>name</code> attribute from the XMLSchema <code>simpleType</code> element representing the renamed type.
<code>type</code>	The XMLSchema type for the base type.
<code>repositoryID</code>	The fully specified repository ID for the CORBA type.

Examples

For example, the definition of `myLong` in [Example 10](#), can be described as shown

Example 10: *myLong IDL*

```
//IDL
typedef long myLong;
```

in [Example 11](#):

Example 11: *myLong WSDL*

```
<?xml version="1.0" encoding="UTF-8"?>
<definitions name="typedef.idl" ...>
```

Example 11: *myLong WSDL*

```

<types>
...
  <xsd:simpleType name="myLong">
    <xsd:restriction base="xsd:int"/>
  </xsd:simpleType>
...
</types>
...
<corba:typeMapping
  targetNamespace="http://schemas.ionac.com/bindings/corba/typeMap"
  <corba:alias name="myLong" type="xsd:int"
    repositoryID="IDL:myLong:1.0" basetype="corba:long"/>
</corba:typeMapping>
</definitions>

```

corba:array**Synopsis**

```

<corba:array name="..." repositoryID="..." type="..."
  elemtype="..." bound="..." />

```

Description

In the CORBA type map, arrays are described using a `corba:array` element.

Attributes

A `corba:array` has the following required attributes:

<code>name</code>	A unique identifier used to reference the CORBA type in the binding.
<code>repositoryID</code>	The fully specified repository ID for the CORBA type.
<code>type</code>	The logical type the structure is mapping.
<code>elemtype</code>	The IDL type of the array's element. This type can be either a primitive type or another complex type that is defined within the type map.
<code>bound</code>	The size of the array.

Examples

For example, consider an array, `myArray`, as defined in [Example 12](#).

Example 12: *myArray IDL*

```

//IDL
typedef long myArray[10];

```

The array `myArray` will have a CORBA type map description similar to the one shown in [Example 13](#).

Example 13: *myArray CORBA type map*

```
<corba:typeMapping targetNamespace="http://schemas.iona.com/bindings/corba/typemap">
  <corba:array name="myArray" repositoryID="IDL:myArray:1.0" type="xsd:long"
    elementType="corba:long" bound="10"/>
</corba:typeMapping>
```

corba:sequence

Synopsis

```
<corba:sequence name="..." repositoryID="..." elementType="..."
  bound="..." />
```

Description

The `corba:sequence` element represents an IDL sequence.

Attributes

A `corba:sequence` has five required attributes.

<code>name</code>	A unique identifier used to reference the CORBA type in the binding.
<code>repositoryID</code>	The fully specified repository ID for the CORBA type.
<code>type</code>	The logical type the structure is mapping.
<code>elementType</code>	The IDL type of the sequence's elements. This type can be either a primitive type or another complex type that is defined within the type map.
<code>bound</code>	The size of the sequence.

Examples

For example, consider the two sequences defined in [Example 14](#), `longSeq` and `charSeq`.

Example 14: *IDL Sequences*

```
\\ IDL
typedef sequence<long> longSeq;
typedef sequence<char, 10> charSeq;
```

The sequences described in [Example 14](#) has a CORBA type map description similar to that shown in [Example 15](#).

Example 15: *CORBA type map for Sequences*

```
<corba:typeMapping targetNamespace="http://schemas.iona.com/bindings/corba/typemap">
  <corba:sequence name="longSeq" repositoryID="IDL:longSeq:1.0" type="xsd1:longSeq"
    elemtype="corba:long" bound="0"/>
  <corba:sequence name="charSeq" repositoryID="IDL:charSeq:1.0" type="xsd1:charSeq"
    elemtype="corba:char" bound="10"/>
</corba:typeMapping>
```

corba:exception

Synopsis

```
<corba:exception name="..." type="..." repositoryID="...">
  <corba:member ... />
  ...
</corba:exception>
```

Description

The `corba:exception` element is a child of a `corba:typeMapping` element. It describes an exception in the CORBA type map. The pieces of data returned with the exception are described by a series of `corba:member` elements. The elements must be declared in the same order as in the IDL representation of the exception.

Attributes

A `corba:exception` element has the following required attributes:

<code>name</code>	A unique identifier used to reference the CORBA type in the binding.
<code>type</code>	The logical type the structure is mapping.
<code>repositoryID</code>	The fully specified repository ID for the CORBA type.

Examples

For example, consider the exception `idNotFound` defined in [Example 16](#).

Example 16: *idNotFound Exception*

```
\\IDL
exception idNotFound
{
  short id;
};
```

In the CORBA type map portion of the contract, `idNotFound` is described by an entry similar to that shown in [Example 17](#):

Example 17: *CORBA Type Map for `idNotFound`*

```
<corba:typeMapping targetNamespace="http://schemas.iona.com/bindings/corba/typemap">
...
  <corba:exception name="idNotFound" type="xsd1:idNotFound" repositoryID="IDL:idNotFound:1.0">
    <corba:member name="id" idltype="corba:short"/>
  </corba:exception>
</corba:typeMapping>
```

corba:anosequence

Synopsis

```
<corba:anosequence name="..." bound="..." elemtype="..."
type="..." />
```

Description

The `corba:anosequence` element is used when representing recursive types. Because XMLSchema recursion requires the use of two defined types and IDL recursion does not, the CORBA type map uses the `corba:anosequence` element as a means of bridging the gap. When Artix generates IDL from a contract, it will not generate new IDL types for XMLSchema types that are used in a `corba:anosequence` element.

Attributes

The `corba:anosequence` element has four required attributes:

<code>name</code>	A unique identifier used to reference the CORBA type in the binding.
<code>bound</code>	The size of the sequence.
<code>elemtype</code>	The name of the CORBA type map element that defines the contents of the sequence.
<code>type</code>	The logical type the element represents.

Examples

[Example 18](#) shows a recursive XMLSchema type, `allAboutMe`, defined using a named type.

Example 18: *Recursive XML Schema Type*

```
<complexType name="allAboutMe">
  <sequence>
    <element name="shoeSize" type="xsd:int"/>
    <element name="mated" type="xsd:boolean"/>
    <element name="conversation" type="tns:moreMe"/>
  </sequence>
</complexType>
<complexType name="moreMe">
  <sequence>
    <element name="item" type="tns:allAboutMe"
      maxOccurs="unbounded"/>
  </sequence>
</complexType>
```

[Example 19](#) shows the how Artix maps the recursive type into the CORBA type map of an Artix contract.

Example 19: *Recursive CORBA Typemap*

```
<corba:anonsequence name="moreMe" bound="0"
  elemtype="ns1:allAboutMe" type="xsd1:moreMe"/>
<corba:struct name="allAboutMe"
  repositoryID="IDL:allAboutMe:1.0"
  type="xsd1:allAboutMe">
  <corba:member name="shoeSize" idltype="corba:long"/>
  <corba:member name="mated" idltype="corba:boolean"/>
  <corba:member name="conversation" idltype="ns1:moreMe"/>
</corba:struct>
```

While the XML in the CORBA typemap does not explicitly retain the recursive nature of recursive XMLSchema types, the IDL generated from the typemap restores the recursion in the IDL type. The IDL generated from the type map in [Example 19](#) defines `allAboutMe` using recursion. [Example 20](#) shows the generated IDL.

Example 20: *IDL for a Recursive Data Type*

```
\\IDL
struct allAboutMe
{
    long shoeSize;
    boolean mated;
    sequence<allAboutMe> conversation;
};
```

corba:anonstring

Synopsis

```
<corba:anonstring name="..." bound="..." type="..." />
```

Description

The `corba:anonstring` element is used to represent instances of anonymous XMLSchema simple types that are derived from `xsd:string`. As with `corba:anonsequence` elements, `corba:anonstring` elements do not result in generated IDL types.

Attributes

`corba:anonstring` elements have three attributes.

<code>name</code>	A unique identifier used to reference the CORBA type in the binding.
<code>bound</code>	The maximum length of the string.
<code>type</code>	The XMLSchema type of the base type. Typically this is <code>xsd:string</code> .

Examples

The complex type, `madAttr`, described in [Example 21](#) contains a member, `style`, that is an instance of an anonymous type derived from `xsd:string`.

Example 21: *madAttr XML Schema*

```
<complexType name="madAttr">
  <sequence>
    <element name="style">
      <simpleType>
        <restriction base="xsd:string">
          <maxLength value="3"/>
        </restriction>
      </simpleType>
    </element>
    <element name="gender" type="xsd:byte"/>
  </sequence>
</complexType>
```

`madAttr` would generate the CORBA typemap shown in [Example 22](#). Notice that `style` is given an IDL type defined by a `corba:anonstring` element.

Example 22: *madAttr CORBA typemap*

```
<corba:typeMapping targetNamespace="http://schemas.iona.com/anonCat/corba/typemap/">
  <corba:struct name="madAttr" repositoryID="IDL:madAttr:1.0" type="xsd1:madAttr">
    <corba:member idltype="ns1:styleType" name="style"/>
    <corba:member idltype="corba:char" name="gender"/>
  </corba:struct>
  <corba:anonstring bound="3" name="styleType" type="xsd:string"/>
</corba:typeMapping>
```

corba:object**Synopsis**

```
<corba:object binding="..." name="..." repositoryID="..."
type="..." />
```

Description

The `corba:object` element is used to represent Artix references in the CORBA type map.

Attributes

corba:object elements have four attributes:

binding	Specifies the binding to which the object refers. If the annotation element is left off the reference declaration in the schema, this attribute will be blank.
name	Specifies the name of the CORBA type. If the annotation element is left off the reference declaration in the schema, this attribute will be Object. If the annotation is used and the binding can be found, this attribute will be set to the name of the interface that the binding represents.
repositoryID	Specifies the repository ID of the generated IDL type. If the annotation element is left off the reference declaration in the schema, this attribute will be set to IDL:omg.org/CORBA/Object/1.0. If the annotation is used and the binding can be found, this attribute will be set to a properly formed repository ID based on the interface name.
type	Specifies the schema type from which the CORBA type is generated. This attribute is always set to references:Reference.

Examples

[Example 23](#) shows an Artix contract fragment that uses Artix references.

Example 23: Reference Sample

```
<?xml version="1.0" encoding="UTF-8" ?>
<definitions name="bankService"
  targetNamespace="http://schemas.myBank.com/bankTypes"
  xmlns="http://schemas.xmlsoap.org/wsdl/"
  xmlns:tns="http://schemas.myBank.com/bankService"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:xsd1="http://schemas.myBank.com/bankTypes"
  xmlns:corba="http://schemas.iona.com/bindings/corba"
  xmlns:corbatm="http://schemas.iona.com/typemap/corba/bank.idl"
  xmlns:references="http://schemas.iona.com/references">
  <types>
    <schema
      targetNamespace="http://schemas.myBank.com/bankTypes"
      xmlns="http://www.w3.org/2001/XMLSchema"
      xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
      <xsd:import schemaLocation="./references.xsd"
        namespace="http://schemas.iona.com/references"/>
```

Example 23: *Reference Sample (Continued)*

```

...
    <xsd:element name="account" type="references:Reference">
      <xsd:annotation>
        <xsd:appinfo>
          corba:binding=AccountCORBABinding
        </xsd:appinfo>
      </xsd:annotation>
    </xsd:element>
  </schema>
</types>
...
<message name="find_accountResponse">
  <part name="return" element="xsd1:account" />
</message>
<message name="create_accountResponse">
  <part name="return" element="xsd1:account" />
</message>
<portType name="Account">
  <operation name="account_id">
    <input message="tns:account_id" name="account_id" />
    <output message="tns:account_idResponse"
      name="account_idResponse" />
  </operation>
  <operation name="balance">
    <input message="tns:balance" name="balance" />
    <output message="tns:balanceResponse"
      name="balanceResponse" />
  </operation>
  <operation name="withdraw">
    <input message="tns:withdraw" name="withdraw" />
    <output message="tns:withdrawResponse"
      name="withdrawResponse" />
    <fault message="tns:InsufficientFundsException"
      name="InsufficientFunds" />
  </operation>
  <operation name="deposit">
    <input message="tns:deposit" name="deposit" />
    <output message="tns:depositResponse"
      name="depositResponse" />
  </operation>
</portType>

```

Example 23: *Reference Sample (Continued)*

```

<portType name="Bank">
  <operation name="find_account">
    <input message="tns:find_account" name="find_account" />
    <output message="tns:find_accountResponse"
      name="find_accountResponse" />
    <fault message="tns:AccountNotFound"
      name="AccountNotFound" />
  </operation>
  <operation name="create_account">
    <input message="tns:create_account" name="create_account" />
    <output message="tns:create_accountResponse"
      name="create_accountResponse" />
    <fault message="tns:AccountAlreadyExistsException"
      name="AccountAlreadyExists" />
  </operation>
</portType>
</definitions>

```

The element named `account` is a reference to the interface defined by the `Account` port type and the `find_account` operation of `Bank` returns an element of type `account`. The annotation element in the definition of `account` specifies the binding, `AccountCORBABinding`, of the interface to which the reference refers.

[Example 24](#) shows the generated CORBA typemap resulting from generating both the `Account` and the `Bank` interfaces into the same contract.

Example 24: *CORBA Typemap with References*

```

<corba:typeMapping
  targetNamespace="http://schemas.myBank.com/bankService/corba/typemap/">
  ...
  <corba:object binding="" name="Object"
    repositoryID="IDL:omg.org/CORBA/Object/1.0" type="references:Reference" />
  <corba:object binding="AccountCORBABinding" name="Account"
    repositoryID="IDL:Account:1.0" type="references:Reference" />
</corba:typeMapping>

```

There are two entries because `wsdltocorba` was run twice on the same file. The first CORBA object is generated from the first pass of `wsdltocorba` to generate the CORBA binding for `Account`. Because `wsdltocorba` could not find the binding specified in the annotation, it generated a generic `Object` reference. The

second CORBA object, `Account`, is generated by the second pass when the binding for `Bank` was generated. On that pass, `ws1dtocorba` could inspect the binding for the `Account` interface and generate a type-specific object reference.

[Example 25](#) shows the IDL generated for the `Bank` interface.

Example 25: *IDL Generated From Artix References*

```
//IDL
...
interface Account
{
    string account_id();
    float balance();
    void withdraw(in float amount)
        raises(::InsufficientFundsException);
    void deposit(in float amount);
};
interface Bank
{
    ::Account find_account(in string account_id)
        raises(::AccountNotFoundException);
    ::Account create_account(in string account_id,
                           in float initial_balance)
        raises(::AccountAlreadyExistsException);
};
```


Tuxedo FML Binding

Artix supports the use of Tuxedo's FML buffers. It uses a set of Artix specific elements placed in the WSDL binding element.

Runtime Compatibility

The Tuxedo FML extension elements are only compatible with the C++ runtime.

Namespace

The WSDL extensions used for the FML binding are defined in the namespace `http://schemas.ionac.com/transportstuxedo`. Add the following namespace declaration to any contracts that use an FML binding:

```
xmlns:tuxedo="http://schemas.ionac.com/transportstuxedo"
```

FML\XMLSchema Support

An FML buffer can only contain the data types listed in [Table 3](#).

Table 3: *FML Type Support*

XML Schema Type	FML Type
xsd:short	short
xsd:unsignedShort	short
xsd:int	long
xsd:unsignedInt	long
xsd:float	float
xsd:double	double
xsd:string	string
xsd:base64Binary	string
xsd:hexBinary	string

Due to FML limitations, support for complex types is limited to `xsd:sequence` and `xsd:all`.

tuxedo:binding

Synopsis

```
<tuxedo:binding />
```

Description

The `tuxedo:binding` element informs Artix that the payload being described is an FML buffer. It is a child of the WSDL `binding` element and has no children.

tuxedo:fieldTable

Synopsis

```
<tuxedo:fieldTable type="...">
  <tuxedo:field ... />
  ...
</tuxedo:fieldTable>
```

Description

The `tuxedo:fieldTable` element contains the mappings between the elements defined in the logical section of the contract and their associated FML `fieldId`.

Attributes

The `tuxedo:fieldTable` element has one required attribute, `type`, that specifies if the FML buffer is an FML16 buffer or an FML32 buffer. [Table 4](#) shows the values of the `type` attribute.

Table 4: *Values of tuxedo:fieldTable Element's type Attribute*

Value	Meaning
FML	The represented FML buffer is a FML16 buffer.
FML32	The represented FML buffer is an FML32 buffer.

tuxedo:field**Synopsis**

```
<tuxedo:field name="..." id="..." />
```

Description

The `tuxedo:field` element defines the association between an element in the logical contract and its corresponding entry in the physical FML buffer. Each element in a message, either a message part or an element in a complex type, must have a corresponding `tuxedo:field` element in the FML binding.

Attributes

The `tuxedo:field` element takes two attributes:

<code>name</code>	The value of the <code>name</code> attribute from the logical message element to which this <code>tuxedo:field</code> element corresponds.
<code>id</code>	The <code>fieldId</code> value of the corresponding element in the generated C++ header defining the FML buffer.

tuxedo:operation**Synopsis**

```
<tuxedo:operaiton />
```

Description

The `tuxedo:operation` element is a child of the WSDL binding's `operation` element. It informs Artix that the messages used by the operation are being passed as FML buffers.

Fixed Binding

The fixed binding supports mapping between XML Schema message definitions and messages formatted in fixed length records.

Runtime Compatibility

The fixed binding's extension elements are only compatible with the C++ runtime.

Namespace

The extensions used to describe fixed record length messages are defined in the namespace `http://schemas.iona.com/bindings/fixed`. Artix tools use the prefix `fixed` to represent the fixed record length extensions. Add the following line to your contract:

```
xmlns:fixed="http://schemas.iona.com/bindings/fixed"
```

`fixed:binding`

Synopsis

```
<fixed:binding justification="..." encoding="..."
    padHexCode="..." />
```

Description

The `fixed:binding` element is a child of the WSDL `binding` element. It specifies that the binding defines a mapping between fixed record length data and the XMLSchema representation of the data.

Attributes

The `fixed:binding` element has three attributes:

<code>justification</code>	Specifies the default justification of the data contained in the messages. Valid values are <code>left</code> and <code>right</code> . Default is <code>left</code> .
<code>encoding</code>	Specifies the codeset used to encode the text data. Valid values are any valid ISO locale or IANA codeset name. Default is <code>UTF-8</code> .
<code>padHexCode</code>	Specifies the hex value of the character used to pad the record.

The settings for the attributes on the `fixed:binding` element become the default settings for all the messages being mapped to the current binding.

fixed:operation**Synopsis**

```
<fixed:operation discriminator="..." />
```

Description

The `fixed:operation` element is a child element of the WSDL `operation` element and specifies that the operation's messages are being mapped to fixed record length data.

Attributes

The `fixed:operation` element has one attribute, `discriminator`, that assigns a unique identifier to the operation. If your service only defines a single operation, you do not need to provide a discriminator. However, if your operation has more than one service, you must define a unique discriminator for each operation in the service. Not doing so will result in unpredictable behavior when the service is deployed.

fixed:body**Synopsis**

```
<fixed:body justification="..." encoding="..." padHexCode="...">
  ...
</fixed:body>
```

Description

The `fixed:body` element is a child element of the `input`, `output`, and `fault` messages being mapped to fixed record length data. It specifies that the message body is mapped to fixed record length data on the wire and describes the exact mapping for the message's parts.

The order in which the message parts are listed in the `fixed:body` element represent the order in which they are placed on the wire. It does not need to correspond to the order in which they are specified in the WSDL `message` element defining the logical message.

The following child elements are used in defining how logical data is mapped to a concrete fixed format message:

- [fixed:field](#) maps message parts defined using a simple type.
- [fixed:sequence](#) maps message parts defined using a `sequence` complex type.

Note: Complex types defined using `all` are not supported by the fixed binding.

- [fixed:choice](#) maps message parts defined using a `choice` complex type.

Attributes

The `fixed:body` element has three attributes:

<code>justification</code>	Specifies how the data in the messages are justified. Valid values are <code>left</code> and <code>right</code> .
<code>encoding</code>	Specifies the codeset used to encode text data. Valid values are any valid ISO locale or IANA codeset name.
<code>padHexCode</code>	Specifies the hex value of the character used to pad the record.

fixed:field

Synopsis

```
<fixed:field name="..." "size="..." format="..."
    justification="..." fixedValue="..." bindingOnly="...">
  <fixed:enumeration ... />
  ...
</fixed:field>
```

Description

The `fixed:field` element is used to map simple data types to a field in a fixed record length message. It is the child of a `fixed:body` element.

Attributes

The `fixed:field` element has the following attributes:

<code>name</code>	Specifies the name of the logical message part that this element represents. It is a required attribute.
-------------------	--

<code>size</code>	Specifies the maximum number of characters in a message part whose base type is <code>xsd:string</code> . Also used to specify the number of characters in the on-wire values used to represent the values of an enumerated type. For more information see “fixed:enumeration” on page 88 .
<code>format</code>	Specifies how non-string data is formatted when it is placed on the wire. For numerical data, formats are entered using <code>#</code> to represent numerical fields and <code>.</code> to represent decimal places. For example <code>##.##</code> would be used to represent <code>12.04</code> . Also can be used for string data that is a date. Date formats use the standard date format syntax. For example, <code>mm/dd/yy</code> would represent dates such as <code>02/23/04</code> and <code>11/02/98</code> .
<code>justification</code>	Specifies the default justification of the data contained in the field. Valid values are <code>left</code> and <code>right</code> . Default is <code>left</code> .
<code>fixedValue</code>	Specifies the value to use for the represented logical message part. The value of <code>fixedValue</code> is always the value placed on the wire for the represented message part. It will override any values set in the application code.
<code>bindingOnly</code>	Specifies if the field appears in the logical definition of the message. The default value is <code>false</code> . When set to <code>true</code> , this attribute signals Artix that it needs to insert a field into the on-wire message that does not appear in the logical message. <code>bindingOnly</code> is used in conjunction with the <code>fixedValue</code> attribute. The <code>fixedValue</code> attribute is used to specify the data to be written into the binding-only field.

Examples

The following examples show different ways of representing data using a `fixed:field` element:

- [String data](#)
- [Numeric data](#)
- [Dates](#)
- [Binding only records](#)

String data

The logical message part, `raverID`, described in [Example 26](#) would be mapped to a `fixed:field` similar to [Example 27](#).

Example 26: Fixed String Message

```
<message name="fixedStringMessage">
  <part name="raverID" type="xsd:string"/>
</message>
```

In order to complete the mapping, you must know the length of the record field and supply it. In this case, the field, `raverID`, can contain no more than twenty characters.

Example 27: Fixed String Mapping

```
<fixed:field name="raverID" size="20"/>
```

Numeric data

If a field contains a 2-digit numeric value with one decimal place, it would be described in the logical part of the contract as an `xsd:float`, as shown in [Example 28](#).

Example 28: Fixed Record Numeric Message

```
<message name="fixedNumberMessage">
  <part name="rageLevel" type="xsd:float"/>
</message>
```

From the logical description of the message, Artix has no way of determining that the value of `rageLevel` is a 2-digit number with one decimal place because the fixed record length binding treats all data as characters. When mapping `rageLevel` in the fixed binding you would specify its `format` with `##.##`, as shown in [Example 29](#). This provides Artix with the metadata needed to properly handle the data.

Example 29: Mapping Numerical Data to a Fixed Binding

```
<fixed:field name="rageLevel" format="##.##"/>
```

Dates

Dates are specified in a similar fashion. For example, the format of the date 12/02/72 is MM/DD/YY. When using the fixed binding it is recommended that dates are described in the logical part of the contract using `xsd:string`. For example, a message containing a date would be described in the logical part of the contract as shown in [Example 30](#).

Example 30: *Fixed Date Message*

```
<message name="fixedDateMessage">
  <part name="goDate" type="xsd:string"/>
</message>
```

If `goDate` is entered using the standard short date format for US English locales, `mm/dd/yyyy`, you would map it to a fixed record field as shown in [Example 31](#).

Example 31: *Fixed Format Date Mapping*

```
<fixed:field name="goDate" format="mm/dd/yyyy" />
```

Binding only records

If you were sending reports that included a fixed expiration date that you did not want exposed to the application, you could create a binding only record called `expDate`. It would be mapped to the fixed field shown in [Example 32](#).

Example 32: *fixedValue Mapping*

```
<fixed:field name="goDate" bindingOnly="true"
  fixedValue="11/11/2112" />
```

fixed:enumeration**Synopsis**

```
<fixed:enumeration value="..." fixedValue="..." />
```

Description

The `fixed:enumeration` element is a child of a `fixed:body` element. It is used to represent the possible values of an enumerated type and define how those values are represented on the wire.

Attributes

The `fixed:enumeration` element has two required attributes:

<code>value</code>	Is the value of the corresponding enumeration value in the logical description of the message part.
<code>fixedValue</code>	Specifies the string value that will be used to represent the logical value on the wire. The length of the string used is determined by the value of the parent <code>fixed:field</code> element's <code>length</code> attribute.

Examples

If you had an enumerated type with the values `FruityTooty`, `Rainbow`, `BerryBomb`, and `OrangeTango` the logical description of the type would be similar to [Example 33](#).

Example 33: *Ice Cream Enumeration*

```
<xs:simpleType name="flavorType">
  <xs:restriction base="xs:string">
    <xs:enumeration value="FruityTooty" />
    <xs:enumeration value="Rainbow" />
    <xs:enumeration value="BerryBomb" />
    <xs:enumeration value="OrangeTango" />
  </xs:restriction>
</xs:simpleType>
```

When you map the enumerated type, you need to know the concrete representation for each of the enumerated values. The concrete representations can be identical to the logical definitions or some other value. The enumerated type in [Example 33](#) could be mapped to the fixed field shown in [Example 34](#). Using this mapping Artix will write OT to the wire for this field if the enumerations value is set to `OrangeTango`.

Example 34: *Fixed Ice Cream Mapping*

```
<fixed:field name="flavor" size="2">
  <fixed:enumeration value="FruityTooty" fixedValue="FT" />
  <fixed:enumeration value="Rainbow" fixedValue="RB" />
  <fixed:enumeration value="BerryBomb" fixedValue="BB" />
  <fixed:enumeration value="OrangeTango" fixedValue="OT" />
</fixed:field>
```

fixed:choice

Synopsis

```
<fixed:choice name="..." discriminatorName="...">
  <fixed:case ... >
    ...
  </fixed:case>
  ...
</fixed:choice>
```

Description

The `fixed:choice` element is a child of a `fixed:body` element. It maps choice complex types to a field in a fixed record length message. The actual values of the choice are defined using `fixed:case` child elements. A `fixed:choice` element must have a `fixed:case` child element for each possible value defined in the choice complex type it represents.

Attributes

The `fixed:choice` element has the following attributes:

<code>name</code>	Specifies the name of the logical message part the choice element is mapping. This attribute is required.
<code>discriminatorName</code>	Specifies the name of a binding-only field that is used as the discriminator for the union. The binding-only field must be defined as part of the parent <code>fixed:body</code> element and must be capable of representing the discriminator.

fixed:case

Synopsis

```
<fixed:case name="..." fixedValue="...">
  ...
</fixed:case>
```

Description

The `fixed:case` element is a child of the `fixed:choice` element. It describes the complete mapping for an element of a choice complex type to a field in a fixed record length message.

To fully describe how the logical data that is represented by a `fixed:case` element is mapped into a field in a fixed record length message, you need to create a mapping for the logical element using children to the `fixed:case` element. The child elements used to map the part's type to the fixed message are

the same as the possible child elements of a `fixed:body` element. `fixed:field` elements describe simple types. `fixed:choice` elements describe choice complex types. `fixed:sequence` elements describe sequence complex types.

Attributes

The `fixed:case` element has the following required attributes:

<code>name</code>	Specifies the value of the <code>name</code> attribute of the corresponding element in the choice complex type being mapped.
<code>fixedValue</code>	Specifies the discriminator value that selects this case. If the parent <code>fixed:choice</code> element has its <code>discriminatorName</code> attribute set, the value must conform to the format specified for that field.

Examples

[Example 35](#) shows an Artix contract fragment mapping a choice complex type to a fixed record length message.

Example 35: Mapping a Union to a Fixed Record Length Message

```
<?xml version="1.0" encoding="UTF-8"?>
<definitions name="fixedMappingsample"
  targetNamespace="http://www.iona.com/FixedService"
  xmlns="http://schemas.xmlsoap.org/wsdl/"
  xmlns:fixed="http://schemas.iona.com/bindings/fixed"
  xmlns:tns="http://www.iona.com/FixedService"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
<types>
  <schema targetNamespace="http://www.iona.com/FixedService"
    xmlns="http://www.w3.org/2001/XMLSchema"
    xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/">
    <xsd:complexType name="unionStationType">
      <xsd:choice>
        <xsd:choice>
          <xsd:element name="train" type="xsd:string"/>
          <xsd:element name="bus" type="xsd:int"/>
          <xsd:element name="cab" type="xsd:int"/>
          <xsd:element name="subway" type="xsd:string"/>
        </xsd:choice>
      </xsd:choice>
    </xsd:complexType>
    ...
  </types>
  <message name="fixedSequence">
    <part name="stationPart" type="tns:unionStationType"/>
  </message>
  <portType name="fixedSequencePortType">
    ...
  </portType>
```

Example 35: *Mapping a Union to a Fixed Record Length Message*

```

<binding name="fixedSequenceBinding"
  type="tns:fixedSequencePortType">
  <fixed:binding/>
  ...
  <fixed:field name="disc" format="##" bindingOnly="true"/>
  <fixed:choice name="stationPart"
    discriminatorName="disc">
    <fixed:case name="train" fixedValue="01">
      <fixed:field name="name" size="20"/>
    </fixed:case>
    <fixed:case name="bus" fixedValue="02">
      <fixed:field name="number" format="###"/>
    </fixed:case>
    <fixed:case name="cab" fixedValue="03">
      <fixed:field name="number" format="###"/>
    </fixed:case>
    <fixed:case name="subway" fixedValue="04">
      <fixed:field name="name" format="10"/>
    </fixed:case>
  </fixed:choice>
  ...
</binding>
...
</definition>

```

fixed:sequence**Synopsis**

```

<fixed:sequence name="..." occurs="..." counterName="...">
  ...
</fixed:field>

```

Description

The `fixed:sequence` element can be a child to a `fixed:body` element, a `fixed:case` element, or another `fixed:sequence` element. It maps a sequence complex type to a field in a fixed record length message.

To fully describe how the complex type that is represented by a `fixed:sequence` element is mapped into a field in a fixed record length message, you need to create a mapping for each of the complex type's elements using children to the `fixed:sequence` element. The child elements used to map the part's type to the fixed message are the same as the possible child elements

of a `fixed:body` element. `fixed:field` elements describe simple types. `fixed:choice` elements describe choice complex types. `fixed:sequence` elements describe sequence complex types.

Attributes

The `fixed:sequence` element has the following attributes:

<code>name</code>	Specifies the value of the <code>name</code> attribute from the corresponding logical complex type. This attribute is required.
<code>occurs</code>	Specifies the number of times this sequence occurs in the message buffer. This value corresponds the value of the <code>maxOccurs</code> attribute of the corresponding logical complex type.
<code>counterName</code>	Specifies the name of the binding-only field that is used to store the actual number of times this sequence occurs in the on-wire message. The corresponding <code>fixed:field</code> element must have enough digits to hold the any whole number up the value of the <code>occurs</code> attribute.

Examples

A structure containing a name, a date, and an ID number would contain three `fixed:field` elements to fully describe the mapping of the data to the fixed record message. [Example 36](#) shows an Artix contract fragment for such a mapping.

Example 36: Mapping a Sequence to a Fixed Record Length Message

```
<?xml version="1.0" encoding="UTF-8"?>
<definitions name="fixedMappingsample"
  targetNamespace="http://www.iona.com/FixedService"
  xmlns="http://schemas.xmlsoap.org/wsdl/"
  xmlns:fixed="http://schemas.iona.com/bindings/fixed"
  xmlns:tns="http://www.iona.com/FixedService"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <types>
    <schema targetNamespace="http://www.iona.com/FixedService"
      xmlns="http://www.w3.org/2001/XMLSchema"
      xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/">
      <xsd:complexType name="person">
        <xsd:sequence>
          <xsd:element name="name" type="xsd:string"/>
          <xsd:element name="date" type="xsd:string"/>
          <xsd:element name="ID" type="xsd:int"/>
        </xsd:sequence>
      </xsd:complexType>
    ...
  </types>
```

Example 36: *Mapping a Sequence to a Fixed Record Length Message*

```
<message name="fixedSequence">
  <part name="personPart" type="tns:person"/>
</message>
<portType name="fixedSequencePortType">
  ...
</portType>
<binding name="fixedSequenceBinding"
  type="tns:fixedSequencePortType">
  <fixed:binding/>
  ...
  <fixed:sequence name="personPart">
    <fixed:field name="name" size="20"/>
    <fixed:field name="date" format="MM/DD/YY"/>
    <fixed:field name="ID" format="#####"/>
  </fixed:sequence>
  ...
</binding>
...
</definition>
```

Tagged Binding

The tagged binding maps between XMLSchema message definitions and self-describing, variable record length messages.

Runtime Compatibility

The tagged binding's extension elements are only compatible with the C++ runtime.

Namespace

The extensions used to describe tagged data bindings are defined in the namespace `http://schemas.ionas.com/bindings/tagged`. Artix tools use the prefix `tagged` to represent the tagged data extensions. Add the following line to the `definitions` element of your contract:

```
xmlns:tagged="http://schemas.ionas.com/bindings/tagged"
```

tagged:binding

Synopsis

```
<tagged:binding selfDescribing="..." fieldSeparator="..."
    fieldNameValueSeparator="..." scopeType="..."
    flattened="..." messageStart="..." messageEnd="..."
    unscopedArrayElement="..." ignoreUnknownElement="..."
    ignoreCase="..." />
```

Description

The `tagged:binding` element specifies that the binding maps logical messages to tagged data messages.

Attributes

The `tagged:binding` element has the following ten attributes:

<code>selfDescribing</code>	Specifies if the message data on the wire includes the field names. Valid values are <code>true</code> or <code>false</code> . If this attribute is set to <code>false</code> , the setting for <code>fieldNameValueSeparator</code> is ignored. This attribute is required.
<code>fieldSeparator</code>	Specifies the delimiter the message uses to separate fields. Valid values include any character that is not a letter or a number. This attribute is required.
<code>fieldNameValueSeparator</code>	Specifies the delimiter used to separate field names from field values in self-describing messages. Valid values include any character that is not a letter or a number.
<code>scopeType</code>	Specifies the scope identifier for complex messages. Supported values are <code>tab(\t)</code> , <code>curlybrace({ data})</code> , and <code>none</code> . The default is <code>tab</code> .
<code>flattened</code>	Specifies if data structures are flattened when they are put on the wire. If <code>selfDescribing</code> is <code>false</code> , then this attribute is automatically set to <code>true</code> .
<code>messageStart</code>	Specifies a special token at the start of a message. It is used when messages that require a special character at the start of a data sequence. Valid values include any character that is not a letter or a number.
<code>messageEnd</code>	Specifies a special token at the end of a message. Valid values include any character that is not a letter or a number.
<code>unscopedArrayElement</code>	Specifies if array elements need to be scoped as children of the array. If set to <code>true</code> arrays take the form <code>echoArray{myArray=2; item=abc; item=def}</code> . If set to <code>false</code> arrays take the form <code>echoArray{myArray=2; {0=abc; 1=def;}}</code> . Default is <code>false</code> .

<code>ignoreUnknownElements</code>	Specifies if Artix ignores undefined element in the message payload. Default is <code>false</code> .
<code>ignoreCase</code>	Specifies if Artix ignores the case with element names in the message payload. Default is <code>false</code> .

The settings for the attributes on these elements become the default settings for all the messages being mapped to the current binding.

tagged:operation

Synopsis

```
<tagged:operation discriminator="..." discriminatorStyle="..." />
```

Description

The `tagged:operation` element is a child element of the WSDL `operation` element. It specifies that the operation's messages are being mapped to a tagged data message.

Attributes

The `tagged:operation` element takes two optional attributes:

<code>discriminator</code>	Specifies a discriminator to be used by the Artix runtime to identify the WSDL operation that will be invoked by the message receiver.
<code>discriminatorStyle</code>	Specifies how the Artix runtime will locate the discriminator as it processes the message. Supported values are <code>msgname</code> , <code>partlist</code> , <code>fieldvalue</code> , and <code>fieldname</code> .

tagged:body

Synopsis

```
<tagged:body>
  ...
</tagged:body>
```

Description

The `tagged:body` element is a child element of the `input`, `output`, and `fault` messages being mapped to a tagged data format. It specifies that the message body is mapped to tagged data on the wire and describes the exact mapping for the message's parts.

The `tagged:body` element will have one or more of the following child elements:

- [tagged:field](#)
- [tagged:sequence](#)
- [tagged:choice](#)

The children describe the detailed mapping of the XMLSchema message to the tagged data to be sent on the wire.

tagged:field

Synopsis

```
<tagged:field name="..." alias="...">
  <tagged:enumeration ... />
  ...
</tagged:field>
```

The `tagged:field` element is a child of a [tagged:body](#) element. It maps simple types and enumerations to a field in a tagged data message. When describing enumerated types a `tagged:field` element will have one or more [tagged:enumeration](#) child elements.

Attributes

The `tagged:field` element has two attributes:

<code>name</code>	A required attribute that must correspond to the name of the logical message part that is being mapped to the tagged data field.
<code>alias</code>	An optional attribute specifying an alias for the field that can be used to identify it on the wire.

tagged:enumeration

Synopsis

```
<tagged:enumeration value="..." />
```

Description

The `tagged:enumeration` element is a child element of a [tagged:field](#) element. It is used to map the value of an enumerated types to a field in a tagged data message.

Parameters

The `tagged:enumeration` element has one required attribute, `value`, that corresponds to the enumeration value as specified in the logical description of the enumerated type.

Examples

If you had an enumerated type, `flavorType`, with the values `FruityTooty`, `Rainbow`, `BerryBomb`, and `OrangeTango` the logical description of the type would be similar to [Example 37](#).

Example 37: Ice Cream Enumeration

```
<xs:simpleType name="flavorType">
  <xs:restriction base="xs:string">
    <xs:enumeration value="FruityTooty" />
    <xs:enumeration value="Rainbow" />
    <xs:enumeration value="BerryBomb" />
    <xs:enumeration value="OrangeTango" />
  </xs:restriction>
</xs:simpleType>
```

`flavorType` would be mapped to a tagged data field as shown in [Example 38](#).

Example 38: Tagged Data Ice Cream Mapping

```
<tagged:field name="flavor">
  <tagged:enumeration value="FruityTooty" />
  <tagged:enumeration value="Rainbow" />
  <tagged:enumeration value="BerryBomb" />
  <tagged:enumeration value="OrangeTango" />
</tagged:field>
```

tagged:sequence

Synopsis

```
<tagged:sequence name="..." alias="..." occurs="...">
  ...
</tagged:sequence>
```

Description

The `tagged:sequence` element is a child of a [tagged:body](#) element, a `tagged:sequence` element, or a [tagged:case](#) element. It maps arrays and sequence complex types to fields in a tagged data message. A `tagged:sequence` element contains one or more children to map the corresponding logical type's parts to fields in a tagged data message. The child elements can be of the following types:

- [tagged:field](#)
- [tagged:sequence](#)
- [tagged:choice](#)

Attributes

The `tagged:sequence` element has three attributes:

<code>name</code>	Specifies the name of the logical message part that is being mapped into the tagged data message. This is a required attribute.
<code>alias</code>	Specifies an alias for the sequence that can be used to identify it on the wire.
<code>occurs</code>	Specifying the number of times the sequence appears. This attribute is used to map arrays.

Examples

A structure containing a name, a date, and an ID number would contain three `tagged:field` elements to fully describe the mapping of the data to the fixed record message. [Example 39](#) shows an Artix contract fragment for such a mapping.

Example 39: *Mapping a Sequence to a Tagged Data Format*

```
<?xml version="1.0" encoding="UTF-8"?>
<definitions name="taggedDataMappingsample"
  targetNamespace="http://www.iona.com/taggedService"
  xmlns="http://schemas.xmlsoap.org/wsdl/"
  xmlns:fixed="http://schemas.iona.com/bindings/tagged"
  xmlns:tns="http://www.iona.com/taggedService"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <types>
    <schema targetNamespace="http://www.iona.com/taggedService"
      xmlns="http://www.w3.org/2001/XMLSchema"
      xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/">
      <xsd:complexType name="person">
        <xsd:sequence>
          <xsd:element name="name" type="xsd:string"/>
          <xsd:element name="date" type="xsd:string"/>
          <xsd:element name="ID" type="xsd:int"/>
        </xsd:sequence>
      </xsd:complexType>
      ...
    </types>
    <message name="taggedSequence">
      <part name="personPart" type="tns:person"/>
    </message>
    <portType name="taggedSequencePortType">
      ...
    </portType>
```

Example 39: Mapping a Sequence to a Tagged Data Format

```
<binding name="taggedSequenceBinding"
  type="tns:taggedSequencePortType">
  <tagged:binding selfDescribing="false" fieldSeparator="pipe"/>
  ...
  <tagged:sequence name="personPart">
    <tagged:field name="name"/>
    <tagged:field name="date"/>
    <tagged:field name="ID"/>
  </tagged:sequence>
  ...
</binding>
...
</definition>
```

tagged:choice

Synopsis

```
<tagged:choice name="..." discriminatorName="..." alias="...">
  <tagged:case ...>
    ...
</tagged:choice>
```

The `tagged:choice` element is a child of a [tagged:body](#) element, a [tagged:sequence](#) element, or a [tagged:case](#) element. It maps unions to a field in a tagged data message. A `tagged:choice` element may contain one or more [tagged:case](#) child elements to map the cases for the union to a field in a tagged data message.

Parameters

The `tagged:choice` element has three attributes:

<code>name</code>	Specifies the name of the logical message part being mapped into the tagged data message. This is a required attribute.
<code>discriminatorName</code>	Specifies the message part used as the discriminator for the union.
<code>alias</code>	Specifies an alias for the union that can be used to identify it on the wire.

tagged:case

Synopsis

```
<tagged:case value="..." />
```

Description

The `tagged:case` element is a child element of a `tagged:choice` element. It describes the complete mapping of a union's individual cases to a field in a tagged data message. A `tagged:case` element must have one child element to describe the mapping of the case's data to a field, or fields, to a tagged data message. Valid child elements are [tagged:field](#), [tagged:sequence](#), and [tagged:choice](#).

Attributes

The `tagged:case` element has one required attribute, `name`, that corresponds to the name of the case element in the union's logical description.

Examples

[Example 40](#) shows an Artix contract fragment mapping a union to a tagged data format.

Example 40: Mapping a Union to a Tagged Data Format

```

<?xml version="1.0" encoding="UTF-8" ?>
<definitions name="fixedMappingsample"
  targetNamespace="http://www.iona.com/tagService"
  xmlns="http://schemas.xmlsoap.org/wsdl/"
  xmlns:fixed="http://schemas.iona.com/bindings/tagged"
  xmlns:tns="http://www.iona.com/tagService"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <types>
    <schema targetNamespace="http://www.iona.com/tagService"
      xmlns="http://www.w3.org/2001/XMLSchema"
      xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/">
      <xsd:complexType name="unionStationType">
        <xsd:choice>
          <xsd:element name="train" type="xsd:string" />
          <xsd:element name="bus" type="xsd:int" />
          <xsd:element name="cab" type="xsd:int" />
          <xsd:element name="subway" type="xsd:string" />
        </xsd:choice>
      </xsd:complexType>
    ...
  </types>
  <message name="tagUnion">
    <part name="stationPart" type="tns:unionStationType" />
  </message>
  <portType name="tagUnionPortType">
    ...
  </portType>

```

Example 40: *Mapping a Union to a Tagged Data Format*

```
<binding name="tagUnionBinding" type="tns:tagUnionPortType">
  <tagged:binding selfDescribing="false"
    fieldSeparator="comma" />
  ...
  <tagged:choice name="stationPart" discriminatorName="disc">
    <tagged:case name="train">
      <tagged:field name="name" />
    </tagged:case>
    <tagged:case name="bus">
      <tagged:field name="number" />
    </tagged:case>
    <tagged:case name="cab">
      <tagged:field name="number" />
    </tagged:case>
    <tagged:case name="subway">
      <tagged:field name="name" />
    </tagged:case>
  </tagged:choice>
  ...
</binding>
...
</definition>
```


TibrvMsg Binding

The Artix TibrvMsg binding elements describe a mapping between XMLSchema messages and the TibrvMsg messages used by Tibco Rendezvous.

Runtime Compatibility

The TibrvMsg binding's extension elements are only compatible with the C++ runtime.

Namespace

The extensions used to describe TibrvMsg bindings are defined in the namespace `http://schemas.ionas.com/transport/tibrv`. Artix tools use the prefix `tibrv` to represent the tagged data extensions. Add the following line to the `definitions` element of your contract:

```
xmlns:tibrv="http://schemas.ionas.com/transport/tibrv"
```

TIBRVMSG to XMLSchema Type Mapping

Table 5 shows how TibrvMsg data types are mapped to XMLSchema types in Artix contracts.

Table 5: *TIBCO to XMLSchema Type Mapping*

TIBRVMSG	XSD
TIBRVMSG_STRING	xsd:string
TIBRVMSG_BOOL	xsd:boolean
TIBRVMSG_I8	xsd:byte
TIBRVMSG_I16	xsd:short
TIBRVMSG_I32	xsd:int
TIBRVMSG_I64	xsd:long
TIBRVMSG_U8	xsd:unsignedByte
TIBRVMSG_U16	xsd:unsignedShort
TIBRVMSG_U32	xsd:unsignedInt
TIBRVMSG_U64	xsd:unsignedLong
TIBRVMSG_F32	xsd:float
TIBRVMSG_F64	xsd:double
TIBRVMSG_STRING	xsd:decimal
TIBRVMSG_DATETIME ^a	xsd:dateTime
TIBRVMSG_OPAQUE	xsd:base64Binary
TIBRVMSG_OPAQUE	xsd:hexBinary
TIBRVMSG_STRING	xsd:QName
TIBRVMSG_STRING	xsd:nonPositiveInteger
TIBRVMSG_STRING	xsd:negativeInteger
TIBRVMSG_STRING	xsd:nonNegativeInteger

Table 5: *TIBCO to XMLSchema Type Mapping*

TIBRVMSG	XSD
TIBRVMSG_STRING	xsd:positiveInteger
TIBRVMSG_STRING	xsd:time
TIBRVMSG_STRING	xsd:date
TIBRVMSG_STRING	xsd:gYearMonth
TIBRVMSG_STRING	xsd:gMonthDay
TIBRVMSG_STRING	xsd:gDay
TIBRVMSG_STRING	xsd:gMonth
TIBRVMSG_STRING	xsd:anyURI
TIBRVMSG_STRING	xsd:token
TIBRVMSG_STRING	xsd:language
TIBRVMSG_STRING	xsd:NMTOKEN
TIBRVMSG_STRING	xsd:Name
TIBRVMSG_STRING	xsd:NCName
TIBRVMSG_STRING	xsd:ID

a. While TIBRVMSG_DATETIME has microsecond precision, `xsd:dateTime` only supports millisecond precision. Therefore, Artix rounds all times to the nearest millisecond.

tibrv:binding

Synopsis

```
<tibrv:binding stringEncoding="..." stringAsOpaque="...">  
  ...  
</tibrv:binding>
```

Description

The `tibrv:binding` element is a child of the WSDL `binding` element. It identifies that the data is to be packed into a `TibrvMsg`. The `tibrv:binding` element can be used to set a default array policy for the `TibrvMsg` generated by the binding by adding a `tibrv:array` child element.

The `tibrv:binding` element can also define binding-only message data by including child elements. The following elements can be a child:

- [tibrv:msg](#)
- [tibrv:field](#)
- [tibrv:context](#)

Any binding-only data defined at the binding level is attached to all messages that use the binding.

Attributes

The `tibrv:binding` element has the following attributes:

<code>stringEncoding</code>	Specifies the character set used in encoding string data included in the message. The default value is <code>utf-8</code> .
<code>stringAsOpaque</code>	Specifies how string data is passed in messages. <code>false</code> , the default value, specifies that strings data is passed as <code>TIBRVMSG_STRING</code> . <code>true</code> specifies that string data is passed as <code>OPAQUE</code> .

tibrv:operation

Synopsis

```
<tibrv:operation>
  ...
</tibrv:operation>
```

Description

The `tibrv:operation` element is a child of a WSDL `operation` element. It signifies that the messages used for this operation are mapped into a `TibrvMsg` and defines any operation specific array policies and data fields.

A `tibrv:operation` element can specify an operation specific array policy by adding a child `tibrv:array` element. This array policy overrides any array policy set at the binding level.

A `tibrv:operation` element can define binding-only message data to be inserted into all `TibrvMsg` messages generated by the operation by adding children to define the data. The following elements are valid children:

- [tibrv:msg](#)
- [tibrv:field](#)
- [tibrv:context](#)

Any binding-only data defined by a `tibrv:operation` element is attached to all messages generated by the operation.

tibrv:input

Synopsis

```
<tibrv:input messageNameFieldPath="..."
            messageNameFieldValue="..."
            stringEncoding="..."
            stringAsOpaque="...">
  ...
</tibrv:input>
```

Description

The `tibrv:input` element is a child of a WSDL `input` element. It defines the exact mapping of the logical input message to the `TibrvMsg` that is used to make requests on a service. When the `tibrv:input` element does not have any children, it signifies that the default XMLSchema message to `TibrvMsg` message mappings are used. If you want to define a custom mapping from the XMLSchema message to the `TibrvMsg` message, want to add context information to the `TibrvMsg` message, or want to add binding only elements to the `TibrvMsg` message, you can add children to the `tibrv:input` element. Valid child elements include:

- [tibrv:msg](#)
- [tibrv:field](#)
- [tibrv:context](#)

A `tibrv:input` element can specify an operation specific array policy by adding a child [tibrv:array](#) element. This array policy overrides any array policy set at the binding level or the operation level.

Attributes

The `tibrv:input` element has the following attributes:

<code>messageNameFieldPath</code>	Specifies the field path that includes the message name. If this attribute is not specified, the first field in the top level message will be used as the message name and given the value <code>IT_BUS_MESSAGE_NAME</code> .
<code>messageNameFieldValue</code>	Specifies the field value that corresponds to the message name. If this attribute is not specified, the value of the WSDL <code>message</code> element's <code>name</code> attribute will be used.
<code>stringEncoding</code>	Specifies the character set used in encoding string data included in the message. This value will override the value set in tibrv:binding .

<code>stringAsOpaque</code>	Specifies how string data is passed in the message. <code>false</code> specifies that strings data is passed as <code>TIBRVMSG_STRING</code> . <code>true</code> specifies that string data is passed as <code>OPAQUE</code> . This value will override the value set in tibrv:binding .
-----------------------------	--

tibrv:output

Synopsis

```
<tibrv:output messageNameFieldPath="..."
  messageNameFieldValue="..."
  stringEncoding="..."
  stringAsOpaque="...">
  ...
</tibrv:output>
```

Description

The `tibrv:output` element is a child of a WSDL `output` element. It defines the exact mapping of the logical output message to the `TibrvMsg` that is used when responding to requests. When the `tibrv:output` element does not have any children, it signifies that the default XMLSchema message to `TibrvMsg` message mappings are used. If you want to define a custom mapping from the XMLSchema message to the `TibrvMsg` message, want to add context information to the `TibrvMsg` message, or want to add binding only elements to the `TibrvMsg` message, you can add children to the `tibrv:output` element. Valid child elements include:

- [tibrv:msg](#)
- [tibrv:field](#)
- [tibrv:context](#)

A `tibrv:output` element can specify an operation specific array policy by adding a child [tibrv:array](#) element. This array policy overrides any array policy set at the binding level or the operation level.

Attributes

The `tibrv:output` element has the following attributes:

<code>messageNameFieldPath</code>	Specifies the field path that includes the message name. If this attribute is not specified, the first field in the top level message will be used as the message name and given the value <code>IT_BUS_MESSAGE_NAME</code> .
<code>messageNameFieldValue</code>	Specifies the field value that corresponds to the message name. If this attribute is not specified, the value of the WSDL <code>message</code> element's <code>name</code> attribute will be used.
<code>stringEncoding</code>	Specifies the character set used in encoding string data included in the message. This value will override the value set in tibrv:binding .
<code>stringAsOpaque</code>	Specifies how string data is passed in the message. <code>false</code> specifies that strings data is passed as <code>TIRBMSG_STRING</code> . <code>true</code> specifies that string data is passed as <code>OPAQUE</code> . This value will override the value set in tibrv:binding .

`tibrv:array`

Synopsis

```
<tibrv:array elementName="..." integralAsSingleField="..."
    loadSize="..." sizeName="..." />
```

Description

The `tibrv:array` element defines how arrays are mapped into elements as a `TibrvMsg` message. The array mapping properties can be set at any level of granularity by making it the child of different `TibrvMsg` binding elements. The array mapping properties at lower levels always override the array mapping properties. For example, the mapping properties defined by a `tibrv:array` element that is the child of a [tibrv:msg](#) element will override the array mapping properties defined by a `tibrv:array` element that is a child of the parent [tibrv:operation](#) element.

Attributes

The array mapping properties are set using the attributes of the `tibrv:array` element. The `tibrv:array` element has the following attributes:

<code>elementName</code>	Specifies an expression that when evaluated will be used as the name of the <code>TibrvMsg</code> field to which array elements are mapped. The default element naming scheme is to concatenate the value of <code>WSDL element element's name</code> attribute with a counter. For information on specifying naming expressions see “Custom array naming expressions” .
<code>integralAsSingleField</code>	Specifies how scalar array data is mapped into <code>TibrvMsgField</code> instances. <code>true</code> , the default, specifies that arrays are mapped into a single <code>TibrvMsgField</code> . <code>false</code> specifies that each member of an array is mapped into a separate <code>TibrvMsgField</code> .
<code>loadSize</code>	Specifies if the number of elements in an array is included in the <code>TibrvMsg</code> . <code>true</code> specifies that the number of elements in the array is added as a <code>TibrvMsgField</code> in the same <code>TibrvMsg</code> as the array. <code>false</code> , the default, specifies that the number of elements in the array is not included in the <code>TibrvMsg</code> .
<code>sizeName</code>	Specifies an expression that when evaluated will be used as the name of the <code>TibrvMsgField</code> to which the size of the array is written. The default naming scheme is to concatenate the value of <code>WSDL element element's name</code> attribute with <code>@size</code> . For information on specifying naming expressions see “Custom array naming expressions” on page 112 .

Custom array naming expressions

When specifying a naming policy for array element names you use a string expression that combines XML properties, strings, and custom naming functions. For example, you could use the expression

```
concat(xml:attr('name'), '_', counter(1,1))
```

to specify that each element in the array `street` is named `street_n`.

Table 6 shows the available functions for use in building array element names.

Table 6: *Functions Used for Specifying TibrvMsg Array Element Names*

Function	Purpose
<code>xml:attr('attribute')</code>	Inserts the value of the named attribute.
<code>concat(item1, item2, ...)</code>	Concatenates all of the elements into a single string.
<code>counter(start, increment)</code>	Adds an increasing numerical value. The counter starts at <i>start</i> and increases by <i>increment</i> .

Examples

Example 41 shows an example of an Artix contract containing a TibrvMsg binding that uses array policies. The policies are set at the binding level and:

- Force the name of the TibrvMsg containing array elements to be named `street0`, `street1`, ...
- Write out the number of elements in each street array.
- Force each element of a street array to be written out as a separate field.

Example 41: *TibrvMsg Binding with Array Policies Set*

```
<?xml version="1.0" encoding="UTF-8"?>
<definitions name="widgetOrderForm.wsdl"
targetNamespace="http://widgetVendor.com/widgetOrderForm"
xmlns="http://schemas.xmlsoap.org/wsdl/"
xmlns:tns="http://widgetVendor.com/widgetOrderForm"
xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns:tibrv="http://schemas.iona.com/transport/tibrv"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:xsd1="http://widgetVendor.com/types/widgetTypes">
  <types>
    <schema
targetNamespace="http://widgetVendor.com/types/widgetTypes"
xmlns="http://www.w3.org/2001/XMLSchema"
xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/">
```

Example 41: *TibrvMsg Binding with Array Policies Set (Continued)*

```

<xsd:complexType name="Address">
  <xsd:sequence>
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="street" type="xsd:string" minOccurs="1"
maxOccurs="5"
                nillable="true"/>
    <xsd:element name="city" type="xsd:string"/>
    <xsd:element name="state" type="xsd:string"/>
    <xsd:element name="zipCode" type="xsd:string"/>
  </xsd:sequence>
</xsd:complexType>
</schema>
</types>
<message name="addressRequest">
  <part name="resident" type="xsd:string"/>
</message>
<message name="addressResponse">
  <part name="address" type="xsd1:Address"/>
</message>
<portType name="theFourOneOne">
  <operation name="lookUp">
    <input message="tns:addressRequest" name="request"/>
    <output message="tns:addressResponse" name="response"/>
  </operation>
</portType>
<binding name="lookUpBinding" type="tns:theFourOneOne">
  <tibrv:binding>
    <tibrv:array elementName="concat(xml:attr('name'),
counter(0, 1))"
                integralsAsSingleField="false"
                loadSize="true"/>
  </tibrv:binding>
  <operation name="lookUp">
    <tibrv:operation/>
    <input name="addressRequest">
      <tibrv:input/>
    </input>
    <output name="addressResponse">
      <tibrv:output/>
    </output>
  </operation>
</binding>

```

Example 41: *TibrvMsg Binding with Array Policies Set (Continued)*

```
<service name="orderWidgetsService">
  <port name="widgetOrderPort" binding="tns:orderWidgetsBinding">
    ...
  </port>
</service>
</definitions>
```

tibrv:msg

Synopsis

```
<tibrv:msg name="..." alias="..." element="..." id="..."
  minOccurs="..." maxOccurs="...">
  ...
</tibrv:msg>
```

Description

The `tibrv:msg` element instructs Artix to create an instance of a `TibrvMsg`.

Attributes

The `tibrv:msg` element has the following attributes:

<code>name</code>	Specifies the name of the contract element which this <code>TibrvMsg</code> instance gets its value. If this attribute is not present, then the <code>TibrvMsg</code> is considered a binding-only element.
<code>alias</code>	Specifies the value of the <code>name</code> member of the <code>TibrvMsg</code> instance. If this attribute is not specified, then the binding will use the value of the <code>name</code> attribute.
<code>element</code>	Used only when <code>tibrv:msg</code> is an immediate child of <code>tibrv:context</code> . Specifies the <code>QName</code> of the element defining the context data to use when populating the <code>TibrvMsg</code> .
<code>id</code>	Specifies the value of the <code>id</code> member of the <code>TibrvMsg</code> instance. The default value is 0.
<code>minOccurs/</code> <code>maxOccurs</code>	Used only with elements that correspond to logical message parts. The values must be identical to the values specified in the schema definition.

tibrv:field

Synopsis

```
<tibrv:field name="..." alias="..." element="..." id="..."
            type="..." value="..." minOccurs="..." maxOccurs="..." />
```

Description

The `tibrv:field` element instructs Artix to create an instance of a `TibrvMsgField`.

Parameters

The `tibrv:field` element has the following attributes:

<code>name</code>	Specifies the name of the contract element which this <code>TibrvMsgField</code> instance gets its value. If this attribute is not present, then the <code>TibrvMsgField</code> is considered a binding-only element.
<code>alias</code>	Specifies the value of the <code>name</code> member of the <code>TibrvMsgField</code> instance. If this attribute is not specified, then the binding will use the value of the <code>name</code> attribute.
<code>element</code>	Used only when <code>tibrv:field</code> is an immediate child of <code>tibrv:context</code> . Specifies the QName of the element defining the context data to use when populating the <code>TibrvMsgField</code> .
<code>id</code>	Specifies the value of the <code>id</code> member of the <code>TibrvMsgField</code> instance. The default value is 0.
<code>type</code>	Specifies the XML Schema type of the data being used to populate the <code>data</code> member of the <code>TibrvMsgField</code> instance.
<code>value</code>	Specifies the value inserted into the <code>data</code> member of the <code>TibrvMsgField</code> instance when the field is a binding-only element.
<code>minOccurs/</code> <code>maxOccurs</code>	Used only with elements that correspond to logical message parts. The values must be identical to the values specified in the schema definition.

tibrv:context

Synopsis

```
<tibrv:context>
  ...
</tibrv:context>
```

Description

The `tibrv:context` element specifies that the following message parts are populated from an Artix context. The child of a `tibrv:context` element can be either:

- a `tibrv:msg` element if the context data is a complex type.
- a `tibrv:msg` element if you wanted to wrap the context data with a `TibrvMsg` on the wire.
- a `tibrv:field` element if the context data is a native XML Schema type.

When a `tibrv:msg` element or a `tibrv:field` element are used to insert context information into a `TibrvMsg` they use the `element` attribute in place of the `name` attribute. The `element` attribute specifies the QName used to register the context data with Artix bus. It must correspond to a globally defined XML Schema element. Also, when inserting context information you cannot specify values for any other attributes except the `alias` attribute.

Examples

If you were integrating with a Tibco server that used a header to correlate messages using an ASCII correlation ID, you could use the `TibrvMsg` binding's context support to implement the correlation ID on the Artix side of the solution. The first step would be to define an XML Schema element called `corrID` for the context that would hold the correlation ID. Then in your `TibrvMsg` binding definition you would include a `tibrv:context` element in the `tibrv:binding` element to specify that all messages passing through the binding will have the header. [Example 42](#) shows a contract fragment containing the appropriate entries for this scenario.

Example 42: *Using Context Data in a TibrvMsg Binding*

```
<definitions
  xmlns:xsd1="http://widgetVendor.com/types/widgetTypes"
  ...>
```

Example 42: *Using Context Data in a TibrvMsg Binding*

```

<types>
  <schema
    targetNamespace="http://widgetVendor.com/types/widgetTypes"
    xmlns="http://www.w3.org/2001/XMLSchema"
    xmlns:wSDL="http://schemas.xmlsoap.org/wsdl/">
    ...
    <element name="corrID" type="xsd:string"/>
    ...
  </schema>
</types>
...
<portType name="correlatedService">
...
</portType>
<binding name="tibrvCorrBinding" type="correlatedService">
  <tibrv:binding>
    <tibrv:context>
      <tibrv:field element="xsd1:corrID"/>
    </tibrv:context>
  </tibrv:binding>
  ...
</binding>
...
</definitions>

```

The context for `corrID` will be registered with the Artix bus using the QName `"http://widgetVendor.com/types/widgetTypes", "corrID"`.

See also

For information on using contexts in Artix applications, see [Developing Artix Applications with C++](#).

XML Binding

Artix includes a binding that supports the exchange of XML documents without the overhead of a SOAP envelope.

Runtime Compatibility

The XML binding's extensions are compatible with the C++ runtime.

Namespace

The extensions used to describe XML format bindings are defined in the namespace `http://celtix.objectweb.org/bindings/xmlformat`. Artix tools use the prefix `xformat` to represent the XML binding extensions. Add the following line to your contracts:

```
xmlns:xformat="http://celtix.objectweb.org/bindings/xmlformat"
```

xformat:binding

Synopsis

```
<xformat:binding rootNode="..." />
```

Description

The `xformat:binding` element is the child of the WSDL `binding` element. It signifies that the messages passing through this binding will be sent as XML documents without a SOAP envelope.

Attributes

The `xformat:binding` element has a single optional attribute called `rootNode`. The `rootNode` attribute specifies the QName for the element that serves as the root node for the XML document generated by Artix. When the `rootNode` attribute is not set, Artix uses the root element of the message part as the root element when using doc style messages or an element using the message part name as the root element when using RCP style messages.

xformat:body

Synopsis

```
<xformat:body rootNode="..." />
```

Description

The `xformat:body` element is an optional child of the WSDL `input` element, the WSDL `output` element, and the WSDL `fault` element. It is used to override the value of the `rootNode` attribute specified in the binding's [xformat:binding](#) element.

Attributes

The `xformat:body` element has a single attribute called `rootNode`. The `rootNode` attribute specifies the QName for the element that serves as the root node for the XML document generated by Artix. When the `rootNode` attribute is not set, Artix uses the root element of the message part as the root element when using doc style messages or an element using the message part name as the root element when using RCP style messages.

Pass Through Binding

The pass through binding is a simple binding that passes blobs through the message layers. The application level code must know how to handle the incoming data.

Runtime Compatibility

The pass through binding's extension elements are only compatible with the C++ runtime.

Namespace

The extensions used to describe tagged data bindings are defined in the namespace `http://schemas.iona.com/bindings/tagged`. Artix tools use the prefix `tagged` to represent the tagged data extensions. Add the following line to the `definitions` element of your contract:

```
xmlns:passthru="http://schemas.iona.com/bindings/passthru"
```

tagged:binding

Synopsis

```
<passthru:binding />
```

Description

The `passthru:binding` element specifies that the binding passes the message through as a blob.

Part II

Ports

In this part

This part contains the following chapters:

HTTP Port	page 125
CORBA Port	page 145
IIOP Tunnel Port	page 149
WebSphere MQ Port	page 153
JMS Port	page 175
Tuxedo Port	page 173
Tibco/Rendezvous Port	page 181
File Transfer Protocol Port	page 191

HTTP Port

Along with the standard WSDL elements used to specify the location of an HTTP port, Artix uses a number of extensions for fine tuning the configuration of an HTTP port.

In this chapter

This chapter discusses the following topics:

Standard WSDL Elements	page 126
Configuration Extensions for C++	page 127
Attribute Details	page 133

Standard WSDL Elements

http:address

Synopsis

```
<http:address location="..." />
```

Description

The `http:address` element is a child of the `WSDL port` element. It specifies the address of the HTTP port of a service that is not using SOAP messages to communicate.

Attributes

The `http:address` element has a single required attribute called `location`. The `location` attribute specifies the service's address as a URL.

soap:address

Synopsis

```
<soap:address location="..." />
```

Description

The `soap:address` element is a child of the `WSDL port` element. It specifies the address of the HTTP port of a service that uses SOAP 1.1 messages to communicate.

Attributes

The `soap:address` element has a single required attribute called `location`. The `location` attribute specifies the service's address as a URL.

wsoap12:address

Synopsis

```
<wsoap12:address location="..." />
```

Description

The `wsoap12:address` element is a child of the `WSDL port` element. It specifies the address of the HTTP port of a service that uses SOAP 1.2 messages to communicate.

Attributes

The `wsoap12:address` element has a single required attribute called `location`. The `location` attribute specifies the service's address as a URL.

Configuration Extensions for C++

Namespace

[Example 43](#) shows the namespace entries you need to add to the `definitions` element of your contract to use the Artix C++ runtime's HTTP extensions.

Example 43: *Artix HTTP Extension Namespaces*

```
<definitions
  ...
  xmlns:http-conf="http://schemas.ionas.com/transport/http/configuration"
  ... >
```

http-conf:client

Synopsis

```
<http-conf:client SendTimeout="..." RecieveTimeout="..."
  AutoRedirect="..." UserName="..."
  Password="..." AuthorizationType="..."
  Authorization="..." Accept="..."
  AcceptLanguage="..." AcceptEncoding="..."
  ContentType="..." Connection="..."
  Host="..." ConnectionAttempts="..."
  CacheControl="..." Cookie="..."
  BrowserType="..." Refferer="..."
  ProxyServer="..." ProxyUsername="..."
  ProxyPassword="..." ProxyAuthorizationType="..."
  ProxyAuthorization="..." UseSecureSockets="..."
  ClientCertificates="..." ClientCertificateChain="..."
  ClientPrivateKey="..." ClientPrivateKeyPassword="..."
  TrustedRootCertificate="..." />
```

Description

The `http-conf:client` element is a child of the WSDL `port` element. It is used to specify client-side configuration details.

Attributes

The `http-conf:client` element has the following attributes:

<code>SendTimeout</code>	Specifies the length of time, in milliseconds, the client tries to send a request to the server before the connection is timed out. Default is 30000.
<code>ReceiveTimeout</code>	Specifies the length of time, in milliseconds, the client tries to receive a response from the server before the connection is timed out. The default is 30000.
<code>AutoRedirect</code>	Specifies if a request should be automatically redirected when the server issues a redirection reply via <code>RedirectURL</code> . The default is <code>false</code> , to let the client redirect the request itself.
<code>UserName</code>	Specifies the user name that the client will use for authentication with a service. This value is passed as an attribute in each request's transport header.
<code>Password</code>	Specifies the password that the client will use for authentication with a service. This value is passed as an attribute in each request's transport header.
<u>AuthorizationType</u>	Specifies the name of the authorization scheme the client wishes to use.
<u>Authorization</u>	Specifies the authorization credentials used to perform the authorization.
<u>Accept</u>	Specifies what media types the client is prepared to handle.
<u>AcceptLanguage</u>	Specifies the client's preferred language for receiving responses.
<u>AcceptEncoding</u>	Specifies what content codings the client is prepared to handle.
<u>ContentType</u>	Specifies the media type of the data being sent in the body of the client request.
<u>Host</u>	Specifies the Internet host and port number of the resource on which the client request is being invoked.

<u>Connection</u>	Specifies if the client wants a particular connection to be kept open after each request/response dialog.
ConnectionAttempts	Specifies the number of times a client will transparently attempt to connect to server.
<u>CacheControl</u>	Specifies directives about the behavior that must be adhered to by caches involved in the chain comprising a request from a client to a server.
Cookie	Specifies a static cookie to be sent to the server along with all requests.
<u>BrowserType</u>	Specifies information about the browser from which the client request originates.
<u>Referer</u>	Specifies the URL of the resource that directed the client to make requests on a particular service.
<u>ProxyServer</u>	Specifies the URL of the proxy server, if one exists along the message path.
ProxyUserName	Specifies the username to use for authentication on the proxy server if it requires separate authorization.
ProxyPassword	Specifies the password to use for authentication on the proxy server if it requires separate authorization.
<u>ProxyAuthorizationType</u>	Specifies the name of the authorization scheme used with the proxy server.
<u>ProxyAuthorization</u>	Specifies the authorization credentials used to perform the authorization with the proxy server.
<u>UseSecureSockets</u>	Indicates if the client wants to open a secure connection.
ClientCertificate	Specifies the full path to the PKCS12-encoded X509 certificate issued by the certificate authority for the client.
ClientCertificateChain	Specifies the full path to the file that contains all the certificates in the chain.

<code>ClientPrivateKey</code>	Specifies the full path to the PKCS12-encoded private key that corresponds to the X509 certificate specified by <code>ClientCertificate</code> .
<code>ClientPrivateKeyPassword</code>	Specifies a password that is used to decrypt the PKCS12-encoded private key.
<code>TrustedRootCertificate</code>	Specifies the full path to the PKCS12-encoded X509 certificate for the certificate authority.

http-conf:server

Synopsis

```
<http_conf:server SendTimeout="..." RecieveTimeout="..."
  SurpressClientSendErrors="..."
  SurpressClientRecieveErrors="..."
  HonnorKeepAlive="..." RedirectURL="..."
  CacheControl="..." ContentLocation="..."
  ContentType="..." ContentEncoding="..."
  ServerType="..." UseSecureSockets="..."
  ServerCertificate="..." ServerCertificateChain="..."
  ServerPrivateKey="..." ServerPrivateKeyPassword="..."
  TrustedRootCertificate="..." />
```

Description

The `http-conf:server` element is a child of the WSDL port element. It is used to specify server-side configuration details.

Attributes

The `http-conf:server` element has the following attributes:

<code>SendTimeout</code>	Sets the length of time, in milliseconds, the server tries to send a response to the client before the connection times out. The default is 30000.
<code>ReceiveTimeout</code>	Sets the length of time, in milliseconds, the server tries to receive a client request before the connection times out. The default is 30000.

<code>SuppressClientSendErrors</code>	Specifies whether exceptions are to be thrown when an error is encountered on receiving a client request. The default is <code>false</code> ; exceptions are thrown on encountering errors.
<code>SuppressClientReceiveErrors</code>	Specifies whether exceptions are to be thrown when an error is encountered on sending a response to a client. The default is <code>false</code> ; exceptions are thrown on encountering errors.
<code>HonorKeepAlive</code>	Specifies whether the server honors client requests for a connection to remain open after a response has been sent. The default is <code>Keep-Alive</code> ; <code>Keep-alive</code> requests are honored. <code>false</code> specifies that <code>keep-alive</code> requests are ignored.
<u>RedirectURL</u>	Sets the URL to which the client request should be redirected if the URL specified in the client request is no longer appropriate for the requested resource.
<u>CacheControl</u>	Specifies directives about the behavior that must be adhered to by caches involved in the chain comprising a response from a server to a client.
<code>ContentLocation</code>	Sets the URL where the resource being sent in a server response is located.
<u>ContentType</u>	Sets the media type of the information being sent in a server response, for example, <code>text/html</code> or <code>image/gif</code> .
<u>ContentEncoding</u>	Specifies what additional content codings have been applied to the information being sent by the server.
<code>ServerType</code>	Specifies what type of server is sending the response to the client. Values take the form <i>program-name/version</i> . For example, <code>Apache/1.2.5</code> .
<u>UseSecureSockets</u>	Indicates whether the server wants a secure HTTP connection running over SSL or TLS.

<code>ServerCertificate</code>	Sets the full path to the PKCS12-encoded X509 certificate issued by the certificate authority for the server.
<u>ServerCertificateChain</u>	Sets the full path to the file that contains all the certificates in the server's certificate chain.
<code>ServerPrivateKey</code>	Sets the full path to the PKCS12-encoded private key that corresponds to the X509 certificate specified by <code>ServerCertificate</code> .
<code>ServerPrivateKeyPassword</code>	Sets a password that is used to decrypt the PKCS12-encoded private key, if it has been encrypted with a password.
<code>TrustedRootCertificate</code>	Sets the full path to the PKCS12-encoded X509 certificate for the certificate authority. This is used to validate the certificate presented by the client.

Attribute Details

AuthorizationType

Description

The `AuthorizationType` attribute corresponds to the HTTP `AuthorizationType` property. It specifies the name of the authorization scheme the client wishes to use. This information is specified and handled at the application level. Artix does not perform any validation on this value. It is the user's responsibility to ensure that the correct scheme name is specified, as appropriate.

Note: If the client wants to use basic username and password-based authentication this does not need to be set.

Authorization

Description

The `Authorization` attribute corresponds to the HTTP `Authorization` property. It specifies the authorization credentials the client wants the server to use when performing the authorization. The credentials are encoded and handled at the application-level. Artix does not perform any validation on the specified value. It is the user's responsibility to ensure that the correct authorization credentials are specified, as appropriate.

Note: If the client wants to use basic username and password-based authentication this does not need to be set.

Accept

Description

The `Accept` attribute corresponds to the HTTP `Accept` property. It specifies what media types the client is prepared to handle. The value of the attribute is specified using as multipurpose internet mail extensions (MIME) types.

MIME type values

MIME types are regulated by the Internet Assigned Numbers Authority (IANA). They consist of a main type and sub-type, separated by a forward slash. For example, a main type of `text` might be qualified as follows: `text/html` or `text/xml`. Similarly, a main type of `image` might be qualified as follows: `image/gif` or `image/jpeg`.

An asterisk (*) can be used as a wildcard to specify a group of related types. For example, if you specify `image/*`, this means that the client can accept any image, regardless of whether it is a GIF or a JPEG, and so on. A value of `*/*` indicates that the client is prepared to handle any type.

Examples of typical types that might be set are:

- `text/xml`
- `text/html`
- `text/text`
- `image/gif`
- `image/jpeg`
- `application/jpeg`
- `application/msword`
- `application/xbitmap`
- `audio/au`
- `audio/wav`
- `video/avi`
- `video/mpeg`

See also

See <http://www.iana.org/assignments/media-types/> for more details.

AcceptLanguage**Description**

The `AcceptLanguage` attribute corresponds to the HTTP `AcceptLanguage` property. It specifies what language (for example, American English) the client prefers for the purposes of receiving a response.

Specifying the language

Language tags are regulated by the International Organization for Standards (ISO) and are typically formed by combining a language code, determined by the ISO-639 standard, and country code, determined by the ISO-3166 standard, separated by a hyphen. For example, `en-US` represents American English.

See also

A full list of language codes is available at <http://www.w3.org/WAI/ER/IG/ert/iso639.htm>.

A full list of country codes is available at

<http://www.iso.ch/iso/en/prods-services/iso3166ma/02iso-3166-code-lists/list-en1.html>.

AcceptEncoding

Description

The `AcceptEncoding` attribute corresponds to the HTTP `AcceptEncoding` Property. It specifies what content encodings the client is prepared to handle. Content encoding labels are regulated by the Internet Assigned Numbers Authority (IANA). Possible content encoding values include `zip`, `gzip`, `compress`, `deflate`, and `identity`.

The primary use of content encodings is to allow documents to be compressed using some encoding mechanism, such as `zip` or `gzip`. Artix performs no validation on content codings. It is the user's responsibility to ensure that a specified content coding is supported at application level.

See also

See <http://www.w3.org/Protocols/rfc2616/rfc2616-sec3.html> for more details on content encodings.

ContentType

Description

The `ContentType` attribute corresponds to the HTTP `ContentType` property. It specifies the media type of the data being sent in the body of a message. Media types are specified using multipurpose internet mail extensions (MIME) types.

MIME type values

MIME types are regulated by the Internet Assigned Numbers Authority (IANA). MIME types consist of a main type and sub-type, separated by a forward slash. For example, a main type of `text` might be qualified as follows: `text/html` or `text/xml`. Similarly, a main type of `image` might be qualified as follows: `image/gif` or `image/jpeg`.

The default type is `text/xml`. Other specifically supported types include:

- `application/jpeg`
- `application/msword`
- `application/xbitmap`
- `audio/au`
- `audio/wav`
- `text/html`
- `text/text`
- `image/gif`
- `image/jpeg`

- video/avi
- video/mpeg.

Any content that does not fit into any type in the preceding list should be specified as `application/octet-stream`.

Client settings

For clients this attribute is only relevant if the client request specifies the `POST` method to send data to the server for processing.

For web services, this should be set to `text/xml`. If the client is sending HTML form data to a CGI script, this should be set to `application/x-www-form-urlencoded`. If the HTTP `POST` request is bound to a fixed payload format (as opposed to SOAP), the content type is typically set to `application/octet-stream`.

See also

See <http://www.iana.org/assignments/media-types/> for more details.

ContentEncoding

Description

The `ContentEncoding` attribute corresponds to the HTTP `ContentEncoding` property. This property specifies any additional content encodings that have been applied to the information being sent by the server. Content encoding labels are regulated by the Internet Assigned Numbers Authority (IANA). Possible content encoding values include `zip`, `gzip`, `compress`, `deflate`, and `identity`.

The primary use of content encodings is to allow documents to be compressed using some encoding mechanism, such as `zip` or `gzip`. Artix performs no validation on content codings. It is the user's responsibility to ensure that a specified content coding is supported at application level.

See also

See <http://www.w3.org/Protocols/rfc2616/rfc2616-sec3.html> for more details on content encodings.

Host

Description

The `Host` attribute corresponds to the HTTP `Host` property. It specifies the internet host and port number of the resource on which the client request is being invoked. This attribute is typically not required. Typically, this attribute does not need to be set. It is only required by certain DNS scenarios or application designs. For example, it indicates what host the client prefers for clusters (that is, for virtual servers mapping to the same internet protocol (IP) address).

Connection

Description

The `Connection` attribute specifies whether a particular connection is to be kept open or closed after each request/response dialog. Valid values are `close` and `Keep-Alive`. The default, `Keep-Alive`, specifies that the client want to keep its connection open after the initial request/response sequence. If the server honors it, the connection is kept open until the client closes it. `close` specifies that the connection to the server is closed after each request/response sequence.

CacheControl

Description

The `CacheControl` attribute specifies directives about the behavior of caches involved in the message chain between clients and servers. The attribute is used for both client and server. However, clients and servers have different settings for specifying cache behavior.

Client-side

[Table 7](#) shows the valid settings for `CacheControl` in `http-conf:client`.

Table 7: *Settings for http-conf:client CacheControl*

Directive	Behavior
<code>no-cache</code>	Caches cannot use a particular response to satisfy subsequent client requests without first revalidating that response with the server. If specific response header fields are specified with this value, the restriction applies only to those header fields within the response. If no response header fields are specified, the restriction applies to the entire response.
<code>no-store</code>	Caches must not store any part of a response or any part of the request that invoked it.
<code>max-age</code>	The client can accept a response whose age is no greater than the specified time in seconds.

Table 7: *Settings for http-conf:client CacheControl*

Directive	Behavior
max-stale	The client can accept a response that has exceeded its expiration time. If a value is assigned to <code>max-stale</code> , it represents the number of seconds beyond the expiration time of a response up to which the client can still accept that response. If no value is assigned, it means the client can accept a stale response of any age.
min-fresh	The client wants a response that will be still be fresh for at least the specified number of seconds indicated.
no-transform	Caches must not modify media type or location of the content in a response between a server and a client.
only-if-cached	Caches should return only responses that are currently stored in the cache, and not responses that need to be reloaded or revalidated.
cache-extension	Specifies additional extensions to the other cache directives. Extensions might be informational or behavioral. An extended directive is specified in the context of a standard directive, so that applications not understanding the extended directive can at least adhere to the behavior mandated by the standard directive.

Server-side

Table 8 shows the valid values for CacheControl in `http-conf:server`.

Table 8: Settings for `http-conf:server` CacheControl

Directive	Behavior
<code>no-cache</code>	Caches cannot use a particular response to satisfy subsequent client requests without first revalidating that response with the server. If specific response header fields are specified with this value, the restriction applies only to those header fields within the response. If no response header fields are specified, the restriction applies to the entire response.
<code>public</code>	Any cache can store the response.
<code>private</code>	Public (<i>shared</i>) caches cannot store the response because the response is intended for a single user. If specific response header fields are specified with this value, the restriction applies only to those header fields within the response. If no response header fields are specified, the restriction applies to the entire response.
<code>no-store</code>	Caches must not store any part of response or any part of the request that invoked it.
<code>no-transform</code>	Caches must not modify the media type or location of the content in a response between a server and a client.
<code>must-revalidate</code>	Caches must revalidate expired entries that relate to a response before that entry can be used in a subsequent response.
<code>proxy-revalidate</code>	Means the same as <code>must-revalidate</code> , except that it can only be enforced on shared caches and is ignored by private unshared caches. If using this directive, the <code>public</code> cache directive must also be used.
<code>max-age</code>	Clients can accept a response whose age is no greater than the specified number of seconds.

Table 8: Settings for `http-conf:server CacheControl` (Continued)

Directive	Behavior
<code>s-maxage</code>	Means the same as <code>max-age</code> , except that it can only be enforced on shared caches and is ignored by private unshared caches. The age specified by <code>s-maxage</code> overrides the age specified by <code>max-age</code> . If using this directive, the <code>proxy-revalidate</code> directive must also be used.
<code>cache-extension</code>	Specifies additional extensions to the other cache directives. Extensions might be informational or behavioral. An extended directive is specified in the context of a standard directive, so that applications not understanding the extended directive can at least adhere to the behavior mandated by the standard directive.

BrowserType

Description

The `BrowserType` attribute specifies information about the browser from which the client request originates. In the HTTP specification from the World Wide Web consortium (W3C) this is also known as the *user-agent*. Some servers optimize based upon the client that is sending the request.

Referer

The `Referer` attribute corresponds to the HTTP Referer property. It specifies the URL of the resource that directed the client to make requests on a particular service. Typically this HTTP property is used when a request is the result of a browser user clicking on a hyperlink rather than typing a URL. This can allow the server to optimize processing based upon previous task flow, and to generate lists of back-links to resources for the purposes of logging, optimized caching, tracing of obsolete or mistyped links, and so on. However, it is typically not used in web services applications.

If the `AutoRedirect` attribute is set to `true` and the client request is redirected, any value specified in the `Referer` attribute is overridden. The value of the HTTP `Referer` property will be set to the URL of the service who redirected the client's original request.

ProxyServer

Description

The `ProxyServer` attribute specifies the URL of the proxy server, if one exists along the message path. A proxy can receive client requests, possibly modify the request in some way, and then forward the request along the chain possibly to the target server. A proxy can act as a special kind of security firewall.

Note: Artix does not support the existence of more than one proxy server along the message path.

ProxyAuthorizationType

Description

The `ProxyAuthorizationType` attribute specifies the name of the authorization scheme the client wants to use with the proxy server. This name is specified and handled at application level. Artix does not perform any validation on this value. It is the user's responsibility to ensure that the correct scheme name is specified, as appropriate.

Note: If basic username and password-based authentication is being used by the proxy server, this does not need to be set.

ProxyAuthorization

Description

The `ProxyAuthorization` attribute specifies the authorization credentials the client will use to perform authorization with the proxy server. These are encoded and handled at application-level. Artix does not perform any validation on the specified value. It is the user's responsibility to ensure that the correct authorization credentials are specified, as appropriate.

Note: If basic username and password-based authentication is being used by the proxy server, this does not need to be set.

UseSecureSockets

Description

The `UseSecureSockets` attribute indicates if the application wants to open a secure connection using SSL or TLS. A secure HTTP connection is commonly referred to as HTTPS. Valid values are `true` and `false`. The default is `false`; the endpoint does not want to open a secure connection.

Note: If the `http:address` element's `location` attribute, or the `soap:address` element's `location` attribute, has a value with a prefix of `https://`, a secure HTTP connection is automatically enabled, even if `UseSecureSockets` is not set to `true`.

RedirectURL

Description

The `RedirectURL` attribute corresponds to the HTTP `RedirectURL` property. It specifies the URL to which the client request should be redirected if the URL specified in the client request is no longer appropriate for the requested resource. In this case, if a status code is not automatically set in the first line of the server response, the status code is set to 302 and the status description is set to `Object Moved`.

ServerCertificateChain

Description

PKCS12-encoded X509 certificates can be issued by intermediate certificate authorities that are not trusted by the client, but which have had their certificates issued in turn by a trusted certificate authority. If this is the case, you can use the `ServerCertificateChain` attribute to allow the certificate chain of PKCS12-encoded X509 certificates to be presented to the client for verification. It specifies the full path to the file that contains all the certificates in the chain.

CORBA Port

Artix supports a robust mechanism for configuring a CORBA endpoint.

Runtime Compatibility

The CORBA transport's extension elements are compatible with the C++ runtime.

C++ Runtime Namespace

The namespace under which the C++ runtime CORBA extensions are defined is `http://schemas.iona.com/bindings/corba`. If you are going to add a C++ runtime CORBA port by hand you will need to add this to your contract's definition element as shown below.

```
xmlns:corba="http://schemas.iona.com/bindings/corba"
```

corba:address

Synopsis

```
<corba:address location="..."/>
```

Description

The `corba:address` element is a child of a WSDL `port` element. It specifies the IOR for the service's CORBA object.

Attributes

The `corba:address` element has one required attribute named `location`. The `location` attribute contains a string specifying the IOR. You have four options for specifying IORs in Artix contracts:

- Entering the object's IOR directly into the contract using the stringified IOR format:

```
IOR:22342...
```

- Entering a file location for the IOR using the following syntax:

```
file:///file_name
```

Note: The file specification requires three backslashes (`///`).

- Entering the object's name using the `corbaname` format:

```
corbaname:rir/NameService#object_name
```

When you use the `corbaname` format for specifying the IOR, Artix will look-up the object's IOR in the CORBA name service.

- Entering the port at which the service exposes itself, using the `corbaloc` syntax.

```
corbaloc:iiop:host:port/service_name
```

corba:policy**Synopsis**

```
<corba:policy poaname="..."|persistent="..."|serviceid="..." />
```

Description

The `corba:policy` element is a child of a WSDL `port` element. It specifies the POA policies the Artix service will use when creating the POA for connecting to a CORBA object. Each `corba:policy` element can only specify one policy. Therefore to define multiple policies you must use multiple `corba:policy` elements.

Attributes

The `corba:policy` element uses attributes to specify the policy it is describing. The following attributes are used:

<code>poaname</code>	Specifies the POA name to use when connecting to the CORBA object. The default POA name is <code>WS_ORB</code> .
<code>persistent</code>	Specifies the value of the POA's persistence policy. The default is <code>false</code> ; the POA is not persistent.
<code>serviceid</code>	Specifies the value of the POA's ID. By default, Artix POAs are assigned their IDs by the ORB.

See also

For more information about CORBA POA policies see the Orbix documentation.

IIOP Tunnel Port

The IIOP tunnel transport allows you to send non-CORBA data over IIOP. This allows you to use a number of the CORBA services.

Runtime Compatibility

The IIOP tunnel transport's extensions are only compatible with the C++ runtime.

Namespace

The namespace under which the IIOP tunnel extensions are defined is `http://schemas.iona.com/bindings/iiop_tunnel`. If you are going to add an IIOP tunnel port by hand you will need to add this to your contract's definition element as shown below.

```
xmlns:iiop="http://schemas.iona.com/bindings/iiop_tunnel"
```

iiop:address

Synopsis

```
<iiop:address location="..."/>
```

Description

The `iiop:address` element is a child of a WSDL `port` element. It specifies the IOR for the CORBA object created for the service.

Attributes

The `iiop:address` element has one required attribute named `location`. The `location` attribute contains a string specifying the IOR. You have four options for specifying IORs in Artix contracts:

- Entering the object's IOR directly into the contract using the stringified IOR format:

```
IOR:22342...
```

- Entering a file location for the IOR using the following syntax:

```
file:///file_name
```

Note: The file specification requires three backslashes (`///`).

- Entering the object's name using the `corbaname` format:

```
corbaname:rir/NameService#object_name
```

When you use the `corbaname` format for specifying the IOR, Artix will look-up the object's IOR in the CORBA name service.

- Entering the port at which the service exposes itself, using the `corbaloc` syntax.

```
corbaloc:iiop:host:port/service_name
```

iiop:payload**Synopsis**

```
<iiop:payload type="..." />
```

Description

The `iiop:payload` element is a child of the WSDL `port` element. It specifies the type of payload being passed through the IIOP tunnel. If the `iiop:payload` element is set, Artix will use the information to attempt codeset negotiation on the contents of the payload being sent through the tunnel. If you do not want codeset negotiation attempted, do not use this element in your IIOP Tunnel port definition.

Attributes

The `iiop:payload` element has a single required element named `type`. The `type` attribute specifies the type of data contained in the payload.

Examples

If your payload contains string data and you want Artix to attempt codeset negotiation you would use the following:

```
<iiop:payload type="string"/>
```

iiop:policy

Synopsis

```
<iiop:policy poaname="..."|persistent="..."|serviceid="..." />
```

Description

The `iiop:policy` element is a child of a WSDL `port` element. It specifies the POA polices the Artix service will use when creating the POA for the IIOP port. Each `iiop:policy` element can only specify one policy. Therefore to define multiple policies you must use multiple `iiop:policy` elements.

Attributes

The `iiop:policy` element uses attributes to specify the policy it is describing. The following attributes are used:

<code>poaname</code>	Specifies the POA name to use when creating the IIOP port. The default POA name is <code>WS_ORB</code> .
<code>persistent</code>	Specifies the value of the POA's persistence policy. The default is <code>false</code> ; the POA is not persistent.
<code>serviceid</code>	Specifies the value of the POA's ID. By default, Artix POAs are assigned their IDs by the ORB.

See also

For more information about CORBA POA policies see the Orbix documentation.

WebSphere MQ Port

Artix provides a number of WSDL extensions to configure a WebSphere MQ service.

In this chapter

This chapter discusses the following topics:

Artix Extension Elements	page 154
Attribute Details	page 160

Artix Extension Elements

Runtime Compatibility

The WebSphere MQ transport's extension elements are only compatible with the C++ runtime.

Namespace

The WSDL extensions used to describe WebSphere MQ transport details are defined in the WSDL namespace `http://schemas.iona.com/transports/mq`. If you are going to use a WebSphere MQ port you need to include the following in the definitions tag of your contract:

```
xmlns:mq="http://schemas.iona.com/transports/mq"
```

mq:client

Synopsis

```
<mq:client QueueManager="..." QueueName="..."
  ReplyQueueManager="..." ReplyQueueName="..."
  Server_Client="..." ModelQueueName="..."
  AliasQueueName="..." ConnectionName="..."
  ConnectionReusable="..." ConnectionFastPath="..."
  UsageStyle="..." CorrelationStyle="..." AccessMode="..."
  Timeout="..." MessageExpiry="..." MessagePriority="..."
  Delivery="..." Transactional="..." ReportOption="..."
  Format="..." MessageID="..." CorrelationID="..."
  ApplicationData="..." AccountingToken="..."
  ApplicationIdData="..." ApplicationOriginData="..."
  UserIdentification="..." />
```

Description

The `mq:client` element is used to configure a client endpoint for connecting to WebSphere MQ. For an MQ client endpoint that receives replies you must provide values for the `QueueManager`, `QueueName`, `ReplyQueueManager`, and `ReplyQueueName` attributes. If the endpoint is not going to receive replies, you do not need to supply settings for the reply queue.

Attributes

The `mq:client` element has the following attributes:

<code>QueueManager</code>	Specifies the name of the queue manager used for making requests.
<code>QueueName</code>	Specifies the name of the queue used for making requests.
<code>ReplyQueueName</code>	Specifies the name of the queue used for receiving responses.
<code>ReplyQueueManager</code>	Specifies the name of the queue manager used for receiving responses.
Server_Client	Specifies which MQ libraries are to be used.
<code>ModelQueueName</code>	Specifies the name of the queue to use as a model for creating dynamic queues.
AliasQueueName	Specifies the local name of the reply queue when the reply queue manager is not on the same host as the client's local queue manager.
<code>ConnectionName</code>	Specifies the name of the connection Artix uses to connect to its queue.
<code>ConnectionReusable</code>	Specifies if the connection can be used by more than one application. The default is <code>false</code> ; the connection is not reusable.
<code>ConnectionFastPath</code>	Specifies if the queue manager will be loaded in process. The default is <code>false</code> ; the queue manager runs as a separate process.
UsageStyle	Specifies if messages can be queued without expecting a response.
CorrelationStyle	Specifies what identifier is used to correlate request and response messages.
AccessMode	Specifies the level of access applications have to the queue.

Timeout	Specifies the amount of time, in milliseconds, between a request and the corresponding reply before an error message is generated.
MessageExpiry	Specifies the value of the MQ message descriptor's <code>Expiry</code> field. It specifies the lifetime of a message in tenths of a second. The default value is <code>INFINITE</code> ; messages never expire.
MessagePriority	Specifies the value of the MQ message descriptor's <code>Priority</code> field.
Delivery	Specifies the value of the MQ message descriptor's <code>Persistence</code> field.
Transactional	Specifies if transaction operations must be performed on the messages.
ReportOption	Specifies the value of the MQ message descriptor's <code>Report</code> field.
Format	Specifies the value of the MQ message descriptor's <code>Format</code> field.
MessageID	Specifies the value of the MQ message descriptor's <code>MsgId</code> field. A value must be specified if <code>CorrelationStyle</code> is set to <code>none</code> .
CorrelationID	Specifies the value for the MQ message descriptor's <code>CorrelId</code> field. A value must be specified if <code>CorrelationStyle</code> is set to <code>none</code> .
ApplicationData	Specifies any application-specific information that needs to be set in the message header.
AccountingToken	Specifies the value for the MQ message descriptor's <code>AccountingToken</code> field.
ApplicationIdData	Specifies the value for the MQ message descriptor's <code>AppIdentityData</code> field.
ApplicationOriginData	Specifies the value for the MQ message descriptor's <code>AppOriginData</code> field.
UserIdentification	Specifies the value for the MQ message descriptor's <code>UserIdentifier</code> field.

mq:server

Synopsis

```
<mq:server QueueManager="..." QueueName="..."
  ReplyQueueManager="..." ReplyQueueName="..."
  Server_Client="..." ModelQueueName="..."
  ConnectionName="..." ConnectionReusable="..."
  ConnectionFastPath="..." UsageStyle="..."
  CorrelationStyle="..." AccessMode="..." Timeout="..."
  MessageExpiry="..." MessagePriority="..." Delivery="..."
  Transactional="..." ReportOption="..." Format="..."
  MessageID="..." CorrelationID="..." ApplicationData="..."
  AccountingToken="..." ApplicationOriginData="..."
  PropagateTransactions="..." />
```

Description

The `mq:server` element is used to configure a server endpoint for connecting to WebSphere MQ. For an MQ server endpoint you must provide values for the `QueueManager` and `QueueName` attributes.

Attributes

The `mq:server` element has the following attributes:

<code>QueueManager</code>	Specifies the name of the queue manager used for receiving requests.
<code>QueueName</code>	Specifies the name of the queue used to receive requests.
<code>ReplyQueueName</code>	Specifies the name of the queue where responses are placed. This setting is ignored if the client specifies a <code>ReplyToQ</code> in a request's message descriptor.
<code>ReplyQueueManager</code>	Specifies the name of the reply queue manager. This setting is ignored if the client specifies a <code>ReplyToQMgr</code> in a request's message descriptor.
<code>Server_Client</code>	Specifies which MQ libraries are to be used.
<code>ModelQueueName</code>	Specifies the name of the queue to use as a model for creating dynamic queues.
<code>ConnectionName</code>	Specifies the name of the connection Artix uses to connect to its queue.

<code>ConnectionReusable</code>	Specifies if the connection can be used by more than one application. The default is <code>false</code> ; the connection is not reusable.
<code>ConnectionFastPath</code>	Specifies if the queue manager will be loaded in process. The default is <code>false</code> ; the queue manager runs as a separate process.
<code>UsageStyle</code>	Specifies if messages can be queued without expecting a response.
<code>CorrelationStyle</code>	Specifies what identifier is used to correlate request and response messages.
<code>AccessMode</code>	Specifies the level of access applications have to the queue.
<code>Timeout</code>	Specifies the amount of time, in milliseconds, between a request and the corresponding reply before an error message is generated.
<code>MessageExpiry</code>	Specifies the value of the MQ message descriptor's <code>Expiry</code> field. It specifies the lifetime of a message in tenths of a second. The default value is <code>INFINITE</code> ; messages never expire.
<code>MessagePriority</code>	Specifies the value of the MQ message descriptor's <code>Priority</code> field.
<code>Delivery</code>	Specifies the value of the MQ message descriptor's <code>Persistence</code> field.
<code>Transactional</code>	Specifies if transaction operations must be performed on the messages.
<code>ReportOption</code>	Specifies the value of the MQ message descriptor's <code>Report</code> field.
<code>Format</code>	Specifies the value of the MQ message descriptor's <code>Format</code> field.
<code>MessageID</code>	Specifies the value of the MQ message descriptor's <code>MsgId</code> field. A value must be specified if <code>CorrelationStyle</code> is set to <code>none</code> .
<code>CorrelationID</code>	Specifies the value for the MQ message descriptor's <code>CorrelId</code> field. A value must be specified if <code>CorrelationStyle</code> is set to <code>none</code> .
<code>ApplicationData</code>	Specifies any application-specific information that needs to be set in the message header.

<code>AccountingToken</code>	Specifies the value for the MQ message descriptor's <code>AccountingToken</code> field.
<code>ApplicationOriginData</code>	Specifies the value for the MQ message descriptor's <code>AppOriginData</code> field.
<code>PropagateTransactions</code>	Specifies if local MQ transactions should be included in flowed transactions. Default is <code>true</code> .

Options

[Table 12](#) describes the correlation between the Artix attribute settings and the `MQOPEN` settings.

Attribute Details

Server_Client

Description

The `Server_Client` attribute specifies which shared libraries to load on systems with a full WebSphere MQ installation.

Parameters

[Table 9](#) describes the settings for this attribute for each type of WebSphere MQ installation.

Table 9: *Server_Client Attribute Settings*

MQ Installation	Server_Client Setting	Behavior
Full		The server shared library (<code>libmqm</code>) is loaded and the application will use queues hosted on the local machine.
Full	<code>server</code>	The server shared library (<code>libmqm</code>) is loaded and the application will use queues hosted on the local machine.
Full	<code>client</code>	The client shared library (<code>libmqic</code>) is loaded and the application will use queues hosted on a remote machine.
Client		The application will attempt to load the server shared library (<code>libmqm</code>) before loading the client shared library (<code>libmqic</code>). The application accesses queues hosted on a remote machine.
Client	<code>server</code>	The application will fail because it cannot load the server shared libraries.
Client	<code>client</code>	The client shared library (<code>libmqic</code>) is loaded and the application accesses queues hosted on a remote machine.

AliasQueueName

Description

The `AliasQueueName` attribute specifies the local name of the reply queue when the service's queue manager is running a different host from the client. Using this attribute ensures that the server will put the replies on the proper queue. Otherwise, the server will receive a request message with the `ReplyToQ` field set to a queue that is managed by a queue manager on a remote host and will be unable to send the reply.

Effect of AliasQueueName

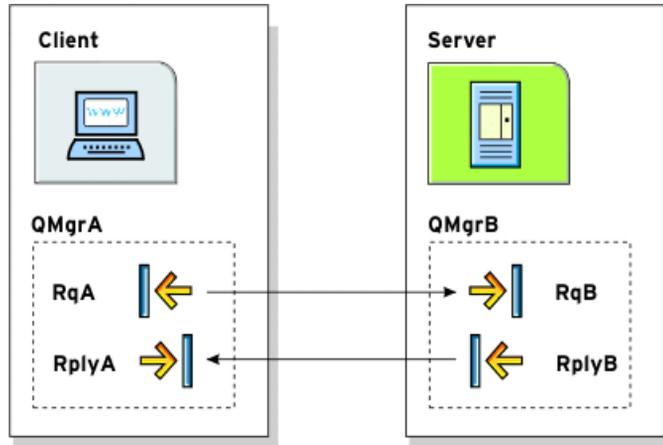
When you specify a value for the `AliasQueueName` attribute in an `mq:client` element, you alter how Artix populates the request's `ReplyToQ` field and `ReplyToQMGr` field. Typically, Artix populates the reply queue information in the request's message descriptor with the values specified in `ReplyQueueManager` and `ReplyQueueName`. Setting `AliasQueueName` causes Artix to leave `ReplytoQMGr` empty and to set `ReplyToQ` to the value of `AliasQueueName`. When the `ReplyToQMGr` field of the message descriptor is left empty, the sending queue manager inspects the queue named in the `ReplyToQ` field to determine who its queue manager is and uses that value for `ReplyToQMGr`. The server puts the message on the remote queue that is configured as a proxy for the client's local reply queue.

Examples

If you had a system defined similar to that shown in [Figure 1](#), you would need to use the `AliasQueueName` attribute setting when configuring your WebSphere MQ client. In this set up the client is running on a host with a local queue manager `QMGrA`. `QMGrA` has two queues configured. `RqA` is a remote queue that is a proxy for `RqB` and `RplyA` is a local queue. The server is running on a different machine whose local queue manager is `QMGrB`. `QMGrB` also has two queues. `RqB` is a local

queue and `RplyB` is a remote queue that is a proxy for `RplyA`. The client places its request on `RqA` and expects replies to arrive on `RplyA`.

Figure 1: *MQ Remote Queues*



The Artix WebSphere MQ port definitions for the client and server for this deployment are shown in [Example 44](#). `AliasQueueName` is set to `RplyB` because that is the remote queue proxying for the reply queue in server's local queue manager. `ReplyQueueManager` and `ReplyQueueName` are set to the client's local queue manager so that it knows where to listen for responses. In this example, the server's `ReplyQueueManager` and `ReplyQueueName` do not need to be set because you are assured that the client is populating the request's message descriptor with the needed information for the server to determine where replies are sent.

Example 44: *Setting Up WebSphere MQ Ports for Intercommunication*

```
<mq:client QueueManager="QMGrA" QueueName="RqA"
  ReplyQueueManager="QMGrA" ReplyQueueName="RplyA"
  AliasQueueName="RplyB"
  Format="string" Convert="true" />
<mq:server QueueManager="QMGrB" QueueName="RqB"
  Format="String" Convert="true" />
```

UsageStyle

Description

The `UsageStyle` specifies if a message can be queued without expecting a response. The default value is `Requester`.

Options

The valid settings for `UsageStyle` are described in [Table 10](#).

Table 10: *UsageStyle Settings*

Attribute Setting	Description
Peer	Specifies that messages can be queued without expecting any response.
Requester	Specifies that the message sender expects a response message. This is the default.
Responder	Specifies that the response message must contain enough information to facilitate correlation of the response with the original message.

Examples

In [Example 45](#), the WebSphere MQ client wants a response from the server and needs to be able to associate the response with the request that generated it. Setting the `UsageStyle` to `responder` ensures that the server's response will properly populate the response message descriptor's `CorrelID` field according to the defined correlation style. In this case, the correlation style is set to `correlationId`.

Example 45: *MQ Client with UsageStyle Set*

```
<mq:client QueueManager="postmaster" QueueName="eddie"
  ReplyQueueManager="postmaster" ReplyQueueName="fred"
  UsageStyle="responder"
  CorrelationStyle="correlationId" />
```

CorrelationStyle

Description

The `CorrelationStyle` attribute specifies how WebSphere MQ matches both the message identifier and the correlation identifier to select a particular message to be retrieved from the queue (this is accomplished by setting the corresponding `MQMO_MATCH_MSG_ID` and `MQMO_MATCH_CORREL_ID` in the `MatchOptions` field in `MQGMO` to indicate that those fields should be used as selection criteria).

Options

The valid correlation styles for an Artix WebSphere MQ port are `messageId`, `correlationId`, and `messageId copy`.

Note: When a value is specified for `ConnectionName`, you cannot use `messageId copy` as the correlation style.

Table 11 shows the actions of `MQGET` and `MQPUT` when receiving a message using a WSDL specified message ID and a WSDL specified correlation ID.

Table 11: *MQGET and MQPUT Actions*

Artix Port Setting	Action for MQGET	Action for MQPUT
<code>messageId</code>	Set the <code>CorrelId</code> of the message descriptor to value of the <code>MessageID</code> .	Copy the value of the <code>MessageID</code> onto the message descriptor's <code>CorrelId</code> .
<code>correlationId</code>	Set <code>CorrelId</code> of the message descriptor to that value of the <code>CorrelationID</code> .	Copy value of the <code>CorrelationID</code> onto message descriptor's <code>CorrelId</code> .
<code>messageId copy</code>	Set <code>MsgId</code> of the message descriptor to value of the <code>messageID</code> .	Copy the value of the <code>MessageID</code> onto message descriptor's <code>MsgId</code> .

AccessMode

Description

The `AccessMode` attribute controls the action of `MQOPEN` and `MQPUT` in the Artix WebSphere MQ transport.

Table 12: *Artix WebSphere MQ Access Modes*

Attribute Setting	Description
peek	peek opens a queue to browse messages. Equivalent to <code>MQOO_BROWSE</code> . This setting is not valid for remote queues.
send	send has the same effect as <code>send+setall</code> for backward compatibility reasons.
send+setall	<p><code>send+setall</code> opens a queue to put messages into. The queue is opened for use with subsequent <code>MQPUT</code> calls. Equivalent to:</p> <pre>MQOPEN => MQOO_SET_ALL_CONTEXT MQOO_OUTPUT MQPUT => MQPMO_SET_ALL_CONTEXT</pre> <p>You can specify different authorizations using other <code>send+</code> settings (for example, <code>send+setid</code>).</p>
send+setid	<p>Equivalent to:</p> <pre>MQOPEN => MQOO_SET_IDENTITY_CONTEXT MQOO_OUTPUT MQPUT => MQPMO_SET_IDENTITY_CONTEXT</pre>
send+passall	<p>Equivalent to:</p> <pre>MQOPEN => MQOO_PASS_ALL_CONTEXT MQOO_OUTPUT MQPUT => MQPMO_PASS_ALL_CONTEXT</pre>
send+passid	<p>Equivalent to:</p> <pre>MQOPEN => MQOO_PASS_IDENTITY_CONTEXT MQOO_OUTPUT MQPUT => MQPMO_PASS_IDENTITY_CONTEXT</pre>
send+none	Equivalent to <code>MQOO_OUTPUT</code> only. This setting has no associated authorization level.

Table 12: *Artix WebSphere MQ Access Modes*

Attribute Setting	Description
receive (default)	receive opens a queue to get messages using a queue-defined default. Equivalent to <code>MQOO_INPUT_AS_Q_DEF</code> . The default value depends on the <code>DefInputOpenOption</code> queue attribute (<code>MQOO_INPUT_EXCLUSIVE</code> or <code>MQOO_INPUT_SHARED</code>).
receive exclusive	receive exclusive opens a queue to get messages with exclusive access. Equivalent to <code>MQOO_INPUT_EXCLUSIVE</code> . The queue is opened for use with subsequent <code>MQGET</code> calls. The call fails with reason code <code>MQRC_OBJECT_IN_USE</code> if the queue is currently open (by this or another application) for input of any type.
receive shared	receive shared opens queue to get messages with shared access. Equivalent to <code>MQOO_INPUT_SHARED</code> . The queue is opened for use with subsequent <code>MQGET</code> calls. The call can succeed if the queue is currently open by this or another application with <code>MQOO_INPUT_SHARED</code> .

MessagePriority

Description

The `MessagePriority` attribute specifies the value for the MQ message descriptor's `Priority` field. Its value must be greater than or equal to zero; zero is the lowest priority. Special values for `MessagePriority` include highest (9), high (7), medium (5), low (3) and lowest (0). The default is normal.

Delivery

Description

The `Delivery` attribute specifies the value of the MQ message descriptor's `Persistence` field.

Options

[Table 13](#) describes the settings for Delviery.

Table 13: *Delivery Attribute Settings*

Artix	WebSphere MQ
persistent	MQPER_PERSISTENT
not persistent (Default)	MQPER_NOT_PERSISTENT

To support transactional messaging, you must make the messages persistent.

Transactional**Description**

The Transactional controls how messages participate in transactions and what role WebSphere MQ plays in the transactions.

Options

The values of the Transactional attribute are explained in [Table 14](#).

Table 14: *Transactional Attribute Settings*

Attribute Setting	Description
none (Default)	The messages are not part of a transaction. No rollback actions will be taken if errors occur.
internal	The messages are part of a transaction with WebSphere MQ serving as the transaction manager.
xa	The messages are part of a flowed transaction with WebSphere MQ serving as an enlisted resource manager.

When the transactional attribute to internal for an Artix service, the following happens during request processing:

1. When a request is placed on the service's request queue, MQ begins a transaction.
2. The service processes the request.
3. Control is returned to the server transport layer.
4. If no reply is required, the local transaction is committed and the request is permanently discarded.

5. If a reply message is required, the local transaction is committed and the request is permanently discarded only after the reply is successfully placed on the reply queue.
6. If an error is encountered while the request is being processed, the local transaction is rolled back and the request is placed back onto the service's request queue.

Examples

Example 46 shows the settings for a WebSphere MQ server port whose requests will be part of transactions managed by WebSphere MQ. Note that the `Delivery` attribute must be set to `persistent` when using transactions.

Example 46: MQ Client Setup to use Transactions

```
<mq:server QueueManager="herman" QueueName="eddie"
  ReplyQueueManager="gomez" ReplyQueueName="lurch"
  UsageStyle="responder" Delivery="persistent"
  CorrelationStyle="correlationId"
  Transactional="internal"/>
```

ReportOption

Description

The `ReportOption` attribute is mapped to the MQ message descriptor's `Report` field. It enables the application sending the original message to specify which report messages are required, whether the application message data is to be included in them, and how the message and correlation identifiers in the report or reply message are to be set. Artix only allows you to specify one `ReportOption` per Artix port. Setting more than one will result in unpredictable behavior.

Options

The values of this attribute are explained in [Table 15](#).

Table 15: *ReportOption Attribute Settings*

Attribute Setting	Description
none (Default)	Corresponds to <code>MQRO_NONE</code> . <code>none</code> specifies that no reports are required. You should never specifically set <code>ReportOption</code> to <code>none</code> ; it will create validation errors in the contract.

Table 15: *ReportOption Attribute Settings*

Attribute Setting	Description
coa	Corresponds to MQRO_COA. coa specifies that confirm-on-arrival reports are required. This type of report is generated by the queue manager that owns the destination queue, when the message is placed on the destination queue.
cod	Corresponds to MQRO_COD. cod specifies that confirm-on-delivery reports are required. This type of report is generated by the queue manager when an application retrieves the message from the destination queue in a way that causes the message to be deleted from the queue.
exception	Corresponds to MQRO_EXCEPTION. exception specifies that exception reports are required. This type of report can be generated by a message channel agent when a message is sent to another queue manager and the message cannot be delivered to the specified destination queue. For example, the destination queue or an intermediate transmission queue might be full, or the message might be too big for the queue.
expiration	Corresponds to MQRO_EXPIRATION. expiration specifies that expiration reports are required. This type of report is generated by the queue manager if the message is discarded prior to delivery to an application because its expiration time has passed.
discard	Corresponds to MQRO_DISCARD_MSG. discard indicates that the message should be discarded if it cannot be delivered to the destination queue. An exception report message is generated if one was requested by the sender

Format

Description

The `Format` attribute is mapped to the MQ message descriptor's `Format` field. It specifies an optional format name to indicate to the receiver the nature of the data in the message.

Options

The value may contain any character in the queue manager's character set, but it is recommended that the name be restricted to the following:

- Uppercase A through Z
- Numeric digits 0 through 9

In addition, the `FormatType` attribute can take the special values `none`, `string`, `event`, `programmable command`, and `unicode`. These settings are described in [Table 16](#).

Table 16: *FormatType Attribute Settings*

Attribute Setting	Description
<code>none</code> (Default)	Corresponds to <code>MQFMT_NONE</code> . No format name is specified.
<code>string</code>	Corresponds to <code>MQFMT_STRING</code> . <code>string</code> specifies that the message consists entirely of character data. The message data may be either single-byte characters or double-byte characters.
<code>unicode</code>	Corresponds to <code>MQFMT_STRING</code> . <code>unicode</code> specifies that the message consists entirely of Unicode characters. (Unicode is not supported in Artix at this time.)
<code>event</code>	Corresponds to <code>MQFMT_EVENT</code> . <code>event</code> specifies that the message reports the occurrence of an WebSphere MQ event. Event messages have the same structure as programmable commands.

Table 16: *FormatType Attribute Settings*

Attribute Setting	Description
programmable command	<p>Corresponds to MQFMT_PCF. programmable command specifies that the messages are user-defined messages that conform to the structure of a programmable command format (PCF) message.</p> <p>For more information, consult the IBM Programmable Command Formats and Administration Interfaces documentation at http://publibfp.boulder.ibm.com/epubs/html/csquizac03/csquizac030d.htm#Header_12.</p>

When you are interoperating with WebSphere MQ applications hosted on a mainframe and the data needs to be converted into the systems native data format, you should set `Format` to `string`. Not doing so will result in the mainframe receiving corrupted data.

Tuxedo Port

Artix can connect to applications that use BEA's Tuxedo as their messaging backbone.

Runtime Compatibility

The Tuxedo transport's extension elements are only compatible with the C++ runtime.

Namespace

The extensions used to describe a Tuxedo port are defined in the namespace `http://schemas.ionas.com/transport/tuxedo`. When a Tuxedo endpoint is defined in a contract, the contract will need the following namespace declaration in the contract's definition element:

```
xmlns:tuxedo="http://schemas.ionas.com/transport/tuxedo"
```

tuxedo:server

Synopsis

```
<tuxedo:server>
  <tuxedo:service ...>
    ...
  </tuxedo:service>
</tuxedo:server>
```

Description

The `tuxedo:server` element is a child of a WSDL `port` element. It contains the definition of a Tuxedo endpoint.

tuxedo:service**Synopsis**

```
<tuxedo:service name="...">
  <tuxedo:input .../>
  ...
</tuxedo:service>
```

Description

The `tuxedo:service` element is the child of a `tuxedo:server` element. It specifies the bulletin board name used to post and receive messages. It has a number of `tuxedo:input` child elements that provide a map to the operations from which messages are routed.

Attributes

The `tuxedo:service` element has a single required attribute called `name`. The `name` attribute specifies the bulletin board name for the service.

tuxedo:input**Synopsis**

```
<tuxedo:input operation="..." />
```

Description

The `tuxedo:input` element specifies which of the operations bound to the port being defined are handled by the Tuxedo service.

Attributes

The `tuxedo:input` element has a single required attribute called `operation`. The `operation` attribute specifies the WSDL operation that is handled by the Tuxedo service. The value must correspond to the value of the `name` attribute of the appropriate WSDL `operation` element.

JMS Port

JMS is a powerful messaging system used by Java applications.

In this chapter

This chapter discusses the following topics:

C++ Runtime Extensions
--

page 176

C++ Runtime Extensions

Namespace

The WSDL extensions used to describe JMS transport details for the C++ runtime are defined in the namespace

`http://celtix.objectweb.org/transport/jms`. If you are going to use a JMS port you need to include the following in the `definitions` tag of your contract:

```
xmlns:jms="http://celtix.objectweb.org/transport/jms"
```

jms:address

Synopsis

```
<jms:address destinationStyle="..."
    jndiConnectionFactoryName="..."
    jndiDestinationName="..."
    jndiReplyDestinationName="..."
    jmsDestinationName="..."
    jmsReplyDestinationName="..."
    connectionUserName="..." connectionPassword="...">
  <jms:JMSNamingProperty ... />
  ...
</jms:address>
```

Description

The `jms:address` element specifies the information needed to connect to a JMS system.

Attributes

The `jms:address` element has the following attributes:

<code>destinationStyle</code>	Specifies if the JMS destination is a JMS queue or a JMS topic.
<code>jndiConnectionFactoryName</code>	Specifies the JNDI name bound to the JMS connection factory to use when connecting to the JMS destination.
<code>jndiDestinationName</code>	Specifies the JNDI name bound to the JMS destination to which Artix connects.

<code>jndiReplyDestinationName</code>	Specifies the JNDI name bound to the JMS destination where replies are sent. This attribute allows you to use a user defined destination for replies.
<code>jmsDestinationName</code>	Specifies the JMS name of the JMS destination to which requests are sent.
<code>jmsReplyDestinationName</code>	Specifies the JMS name of the JMS destination where replies are sent. This attribute allows you to use a user defined destination for replies.
<code>connectionUserName</code>	Specifies the username to use when connecting to a JMS broker.
<code>connectionPassword</code>	Specifies the password to use when connecting to a JMS broker.

jms:JMSNamingProperty

Synopsis

```
<jms:JMSNamingProperty name="..." value="..." />
```

Description

The `jms:JMSNamingProperty` element is a child of the `jms:address` element. It is used to provide the values used to populate the properties object used when connecting to a JNDI provider.

Attributes

The `jms:JMSNamingProperty` element has the following attributes:

<code>name</code>	Specifies the name of the JNDI property to set.
<code>value</code>	Specifies the value for the specified property.

JNDI property names

The following is a list of common JNDI properties that can be set:

- `java.naming.factory.initial`
- `java.naming.provider.url`
- `java.naming.factory.object`
- `java.naming.factory.state`
- `java.naming.factory.url.pkgs`
- `java.naming.dns.url`
- `java.naming.authoritative`
- `java.naming.batchsize`
- `java.naming.referral`
- `java.naming.security.protocol`
- `java.naming.security.authentication`

- `java.naming.security.principal`
- `java.naming.security.credentials`
- `java.naming.language`
- `java.naming.applet`

For more details on what information to use in these attributes, check your JNDI provider's documentation and consult the Java API reference material.

jms:client

Synopsis

```
<jms:client messageType="..." />
```

Description

The `jms:client` element is a child of the `WSDL:port` element. It is used to specify the types of messages being used by a JMS client endpoint and the timeout value for a JMS client endpoint.

Attributes

The `jms:client` element has the following attributes:

<code>messageType</code>	Specifies how the message data will be packaged as a JMS message. <code>text</code> specifies that the data will be packaged as a <code>TextMessage</code> . <code>binary</code> specifies that the data will be packaged as an <code>ObjectMessage</code> .
--------------------------	--

jms:server

Synopsis

```
<jms:server useMessageIDAsCorrelationID="..."
    durableSubscriberName="..."
    messageSelector="..." transactional="..." />
```

Description

The `jms:server` element is a child of the `WSDL:port` element. It specifies settings used to configure the behavior of a JMS service endpoint.

Attributes

The `jms:server` element has the following attributes:

<code>useMessageIDAsCorrealationID</code>	Specifies whether JMS will use the message ID to correlate messages. The default is <code>false</code> .
<code>durableSubscriberName</code>	Specifies the name used to register a durable subscription.
<code>messageSelector</code>	Specifies the string value of a message selector to use.
<code>transactional</code>	Specifies whether the local JMS broker will create transactions around message processing. The default is <code>false</code> .

Tibco/Rendezvous Port

Artix provides a number of attributes to define a TIB/RV service.

In this chapter

This chapter discusses the following topics:

Artix Extension Elements	page 182
Attribute Details	page 187

Artix Extension Elements

Runtime Compatibility

The Tibco/Rendezvous transport's extensions are only compatible with the C++ runtime.

Namespace

The extensions used to describe a Tibco/Rendezvous endpoint are defined in the namespace `http://schemas.ionas.com/transport/tibrv`. When a Tibco endpoint is defined in a contract, the contract will need the following namespace declaration in the contract's `definition` element:

```
xmlns:tibrv="http://schemas.ionas.com/transport/tibrv"
```

tibrv:port

Synopsis

```
<tibrv:port serverSubject="..." clientSubject="..."
  bindingType="..." callbackLevel="..."
  responseDispatchTimeout="..." transportService="..."
  transportNetwork="..." transportDaemon="..."
  transportBatchMode="..." cmSupport="..."
  cmTransportServerName="..." cmTransportClientName="..."
  cmTransportRequestOld="..." cmTransportLedgerName="..."
  cmTransportSyncLedger="..." cmTransportRelayAgent="..."
  cmTransportDefaultTimeLimit="..."
  cmListenerCancelAgreement="..."
  cmQueueTransportServerName="..."
  cmQueueTransportWorkerWeight="..."
  cmQueueTransportWorkerTasks="..."
  cmQueueTransportSchedulerWeight="..."
  cmQueueTransportSchedulerHeartbeat="..."
  cmQueueTransportSchedulerActivation="..."
  cmQueueTransportCompleteTime="..." />
```

Description

The `tibrv:port` element is the child of a WSDL `port` element. It specifies the properties used to configure an endpoint that use Tibco/Rendezvous as its messaging backbone. The element's attributes specify the information needed to configure the transport layer. The `serverSubject` attribute is required to be set and its value must match on both the server side and the client side.

Attributes

The `tibrv:port` element has the following attributes:

<code>serverSubject</code>	Specifies the subject to which the server listens. This parameter must be the same between client and server.
<code>clientSubject</code>	Specifies the prefix to the subject that the client listens to. The default is to use a uniquely generated name.
<code>bindingType</code>	Specifies the message binding type.

<code>callbackLevel</code>	Specifies the server-side callback level when TIB/RV system advisory messages are received.
<code>responseDispatchTimeout</code>	Specifies the client-side response timeout.
<code>transportService</code>	Specifies the UDP service name or port for <code>TibrvNetTransport</code> .
<code>transportNetwork</code>	Specifies the binding network addresses for <code>TibrvNetTransport</code> .
<code>transportDaemon</code>	Specifies the TCP daemon port for <code>TibrvNetTransport</code> . The default is to use 7500 for the <code>TRDP</code> daemon, or 7550 for the <code>PGM</code> daemon.
<code>transportBatchMode</code>	Specifies if the TIB/RV transport uses batch mode to send messages. The default is <code>false</code> ; The endpoint will send messages as soon as they are ready.
<code>cmSupport</code>	Specifies if Certified Message Delivery support is enabled. The default is <code>false</code> ; CM support is disabled.
<code>cmTransportServerName</code>	Specifies the server's <code>TibrvCmTransport</code> correspondent name.
<code>cmTransportClientName</code>	Specifies the client <code>TibrvCmTransport</code> correspondent name. The default is to use a transient correspondent name.
<code>cmTransportRequestOld</code>	Specifies if the endpoint can request old messages on start-up. The default is <code>false</code> ; the endpoint cannot request old messages on start-up.
<code>cmTransportLedgerName</code>	Specifies the <code>TibrvCmTransport</code> ledger file. The default is to use an in-process ledger that is stored in memory.

<code>cmTransportSyncLedger</code>	Specifies if the endpoint uses a synchronous ledger. The default is <code>false</code> ; the endpoint does not use a synchronous ledger.
<code>cmTransportRelayAgent</code>	Specifies the endpoint's <code>TibrvCmTransport</code> relay agent. If this attribute is not set, the endpoint does not use a relay agent.
<code>cmTransportDefaultTimeLimit</code>	Specifies the default time limit for a Certified Message to be delivered. The default is no time limit.
<code>cmListenerCancelAgreements</code>	Specifies if Certified Message agreements are canceled when the endpoint disconnects. The default is <code>false</code> ; agreements remain in place after disconnecting.
<code>cmQueueTransportServerName</code>	Specifies the server's <code>TibrvCmQueueTransport</code> correspondent name.
<code>cmQueueTransportWorkerWeight</code>	Specifies the endpoint's <code>TibrvCmQueueTransport</code> worker weight. The default is <code>TIBRVCM_DEFAULT_WORKER_WEIGHT</code> .
<code>cmQueueTransportWorkerTasks</code>	Specifies the value of the endpoint's <code>TibrvCmQueueTransport</code> worker tasks parameter. The default is <code>TIBRVCM_DEFAULT_WORKER_TASKS</code> .
<code>cmQueueTransportSchedulerWeight</code>	Specifies the value of the <code>TibrvCmQueueTransport</code> scheduler weight parameter. The default is <code>TIBRVCM_DEFAULT_SCHEDULER_WEIGHT</code> .
<code>cmQueueTransportSchedulerHeartbeat</code>	Specifies the value of the <code>TibrvCmQueueTransport</code> scheduler heartbeat parameter. The default is <code>TIBRVCM_DEFAULT_SCHEDULER_HB</code> .

<code>cmQueueTransportSchedulerActivation</code>	Specifies the value of the <code>TibrvCmQueueTransport scheduler activation</code> parameter. The default is <code>TIBRVCM_DEFAULT_SCHEDULER_ACTIVE</code> .
<code>cmQueueTransportCompleteTime</code>	Specifies the value of the <code>TibrvCmQueueTransport complete time</code> parameter. The default is 0.

Attribute Details

bindingType

Description

The `bindingType` attribute specifies the message binding type.

Options

Artix TIB/RV ports support three types of payload formats as described in [Table 17](#).

Table 17: *TIB/RV Supported Payload formats*

Value	Payload Formats	TIB/RV Message Implications
msg	TibrvMsg	The message data is encapsulated in a TibrvMsg described by the binding section of the service's contract.
xml	SOAP, tagged data	The message data is encapsulated in a field of <code>TIBRVMSG_XML</code> with a null name and an ID of 0.
opaque	fixed record length data, variable record length data	The message data is encapsulated in a field of <code>TIBRVMSG_OPAQUE</code> with a null name and an ID of 0.

callbackLevel

Description

The `callbackLevel` attribute specifies the server-side callback level when TIB/RV system advisory messages are received.

Options

It has three settings:

- INFO
- WARN
- ERROR (default)

responseDispatchTimeout

Description

The `responseDispatchTimeout` attribute specifies the client-side response receive dispatch timeout. The default is `TIBRV_WAIT_FOREVER`.

Note: If only the `TibrvNetTransport` is used and there is no server return response for a request, then not setting a timeout value causes the client to block forever.

transportService

Description

The `transportService` attribute specifies the UDP service name or port for `TibrvNetTransport`. The default is `rendezvous`. If no corresponding entry exists in `/etc/services`, 7500 for the `TRDP` daemon, or 7550 for the `PGM` daemon will be used. This parameter must be the same for both client and server.

transportNetwork

Description

The `transportNetwork` attribute specifies the binding network addresses for `TibrvNetTransport`. The default is to use the interface IP address of the host for the `TRDP` daemon, `224.0.1.78` for the `PGM` daemon. This parameter must be interoperable between the client and the server.

cmTransportServerName

Description

The `cmTransportServerName` attribute specifies the server's `TibrvCmTransport` correspondent name. The default is to use a transient correspondent name. This parameter must be the same for both client and server if the client also uses Certified Message Delivery.

cmQueueTransportServerName

Description

The `cmQueueTransportServerName` attribute specifies the server's `TibrvCmQueueTransport` correspondent name. If this property is set, the server listener joins to the distributed queue of the specified name. This parameter must be the same among the server queue members.

File Transfer Protocol Port

Artix can use an FTP server as a middle-tier message broker.

Runtime Compatibility

The FTP transport's extensions are compatible with the C++ runtime.

Namespace

The extensions used to describe a File Transfer Protocol (FTP) port are defined in the namespace `http://schemas.ionas.com/transport/ftp`. When an FTP endpoint is defined in a contract, the contract will need the following namespace declaration in the contract's definition element:

```
xmlns:ftp="http://schemas.ionas.com/transport/ftp"
```

ftp:port

Synopsis

```
<ftp:port host="..." port="..." requestLocation="..."
  replyLocation="..." connectMode="..." scanInterval="...">
  <ftp:properties>
    ...
  </ftp:properties>
</ftp:port>
```

Description

The `ftp:port` element is a child of a WSDL `port` element. It defines the connection details for an FTP endpoint. It may contain an [ftp:properties](#) element.

Attributes

The `ftp:port` element has the following attributes:

<code>host</code>	Specifies the domain name or IP address of the machine hosting the FTPD used by the endpoint.
<code>port</code>	Specifies the port number on which the endpoint will contact the FTPD.
<code>requestLocation</code>	Specifies the path on the FTPD host the endpoint will use for requests. The default is <code>/</code> .
<code>replyLocation</code>	Specifies the path on the FTPD host the endpoint will use for replies. The default is <code>/</code> .
<code>connectMode</code>	Specifies the connection mode used to connect to the FTPD. Valid values are <code>passive</code> and <code>active</code> . The default is <code>passive</code> .
<code>scanInterval</code>	Specifies the interval, in seconds, at which the request and reply directories are scanned for updates. The default is 5.

ftp:properties

Synopsis

```
<ftp:properties>
  <ftp:property ... />
  ...
</ftp:properties>
```

Description

The `ftp:properties` element defines a number of file naming properties used by the endpoint for storing requests and replies. It contains one or more [ftp:property](#) elements.

ftp:property

Synopsis

```
<ftp:property name="..." value="..." />
```

Description

The `ftp:property` element defines specific file naming properties to use when reading and writing messages on the FTPD host. The properties are defined by the implementation used for the naming scheme classes. Artix provides a default implementation. However, a custom naming scheme implementation may have different properties.

Attributes

The `ftp:property` element has the following attributes:

<code>name</code>	Specifies the name of the property to set.
<code>value</code>	Specifies the value of the property.

Default Naming Properties

The default naming implementation provided with Artix supports the following properties:

<code>staticFileNames</code>	Determines if the endpoint uses a static, non-unique, naming scheme for its files. Valid values are <code>true</code> and <code>false</code> . The default is <code>true</code> .
<code>requestFilenamePrefix</code>	Specifies the prefix to use for file names when <code>staticFileNames</code> is set to <code>false</code> .

Part III

Other Extensions

In this part

This part contains the following chapters:

Routing	page 197
Security	page 207
Codeset Conversion	page 211

Routing

Artix provides a number of WSDL extensions for defining how messages are routed between services.

Runtime Compatibility

The extensions described below are only recognized by the Artix router.

Namespace

The Artix routing elements are defined in the `http://schemas.iona.com/routing` namespace. When describing routes in an Artix contract your contract's `definitions` element must have the following entry:

```
xmlns:routing="http://schemas.iona.com/routing"
```

routing:expression

Synopsis

```
<routing:expression name="..." evaluator="..."
  ...
</routing:expression>
```

Description

The `routing:expression` element is a child of the `WSDL definitions` element. It specifies an XPATH expression that evaluates messages for content-based routing.

Attributes

The `routing:expression` requires the following two attributes:

<code>name</code>	Specifies a string that is used to refer to the expression when defining routes.
<code>evaluator</code>	Specifies the name of the grammar used in the expression. Currently the only valid value is <code>xpath</code> .

routing:route**Synopsis**

```
<routing:route name="..." multiRoute="...">
  ...
</routing:route>
```

Description

The `routing:route` element is the root element of each route described in a contract.

Attributes

The `routing:route` element takes the following attributes:

<code>name</code>	Specifies a unique identifier for the route. This attribute is required.
<code>multiRoute</code>	An optional attribute that specifies how messages are sent to the listed destinations. Values are <code>fanout</code> , <code>failover</code> , or <code>loadBalance</code> . Default is to route messages to a single destination.

Options

Standard routes define a single source/destination pair. When the `multiRoute` attribute is specified, your route description will contain more than one destination.

Setting the `multiRoute` attribute has the following effects:

- `fanout` instructs Artix to send messages from the source to all the listed destinations.
 - `failover` instructs Artix to move through the list of destinations until it can successfully send the message.
 - `loadBalance` instructs Artix to use a round-robin algorithm to spread messages across all of the listed destinations.
-

routing:source**Synopsis**

```
<routing:source service="..." port="..." />
```

Description

The `routing:source` element is a child of a [routing:route](#) element. It specifies the port from which the route will redirect messages. A route can have several source elements as long as they all meet the compatibility rules for port-based routing.

Attributes

The `routing:source` element requires two attributes:

<code>service</code>	Specifies the WSDL <code>service</code> element in which the source port is defined.
<code>port</code>	Specifies the name of the WSDL <code>port</code> element from which messages are being received. The router will create a proxy to listen for messages on this port.

routing:query

Synopsis

```
<routing:query expression="...">
  <routing:desitination id="..." ... />
  ...
</routing:query>
```

Description

The `routing:query` element is a child of a [routing:route](#) element. It specifies the destinations for a content-based route. The child [routing:destination](#) elements must use the `id` attribute to specify the value used to select the destination.

Attributes

The `routing:query` element has one attribute:

<code>expression</code>	Specifies the value of the <code>name</code> attribute from the routing:expression element defining the XPATH expression used to select the destination of the message. The query selects the destination with the <code>id</code> value that matches the result of applying the expression to the message content.
-------------------------	---

routing:destination

Synopsis

```
<routing:destination value="..." service="..."
  port="..." route="..." />
```

Description

The `routing:destination` element is a child of a [routing:route](#) element. It specifies the port to which the source messages are directed. The destination must be compatible with all of the source elements.

Attributes

The `routing:destination` element has the following attributes:

<code>value</code>	Specifies the value of the content-based routing query that triggers the destination. This attribute is required when the element is the child of a <code>routing:query</code> element and ignored otherwise.
<code>service</code>	Specifies the WSDL <code>service</code> element in which the destination port is defined.
<code>port</code>	Specifies the name of the port WSDL element to which messages are routed.
<code>route</code>	Specifies a linked route to use for selecting the ultimate destination. When this attribute is used, you should not use the <code>service</code> attribute or the <code>port</code> attribute.

routing:transportAttribute**Synopsis**

```
<routing:transportAttribute>
  ...
</routing:transportAttribute>
```

Description

The `routing:transportAttribute` element is a child of a `routing:route` element. It defines routing rules based on the transport attributes set in a message's header when using HTTP, CORBA, or WebSphere MQ. The criteria for determining if a message meets the transport attribute rule are specified using the following child elements:

- `routing:equals`
- `routing:greater`
- `routing:less`
- `routing:startswith`
- `routing:endswith`
- `routing:contains`
- `routing:empty`
- `routing:nonempty`

A message passes the rule if it meets each criterion specified by the child elements.

Transport attribute rules are defined after all of the operation-based routing rules and before any destinations are listed.

Examples

[Example 47](#) shows a route using transport attribute rules based on HTTP header attributes. Only messages sent to the server whose `UserName` is equal to `JohnQ` will be passed through to the destination port.

Example 47: Transport Attribute Rules

```
<routing:route name="httpTransportRoute">
  <routing:source service="tns:httpService"
    port="tns:httpPort" />
  <routing:transportAttributes>
    <routing:equals
      contextName="http-conf:HTTPServerIncomingContexts"
      contextAttributeName="UserName"
      value="JohnQ" />
    </routing:transportAttributes>
  <routing:destination service="tns:httpDest"
    port="tns:httpDestPort" />
</routing:route>
```

routing:equals

Synopsis

```
<routing:equals contextName="..."
  contextAttributeName="..."
  value="..."
  ignorecase="..." />
```

Description

The `routing:equals` element is a child of a [routing:transportAttribute](#) element. It defines a rule that is triggered when the specified attribute equals the value given. It applies to string or numeric attributes.

Attributes

The `routing:equals` element has the following attributes:

<code>contextName</code>	Specifies the QName of the context in which the desired transport attributes are stored.
<code>contextAttributeName</code>	Specifies the QName of the transport attribute the rule evaluates.
<code>value</code>	Specifies the value against which the specified attribute is evaluated.
<code>ignorecase</code>	Specifies whether the case of characters in a string are ignored. The default is <code>no</code> ; case is considered when evaluating string data.

routing:greater

Synopsis

```
<routing:greater contextName="..."
                contextAttributeName="..."
                value="..." />
```

Description

The `routing:greater` element is a child of a `routing:transportAttribute` element. It defines a rule that is triggered when the value of the specified attribute is greater than the value given. It applies to numeric attributes.

Attributes

The `routing:greater` element has the following attributes:

<code>contextName</code>	Specifies the QName of the context in which the desired transport attributes are stored.
<code>contextAttributeName</code>	Specifies the QName of the transport attribute the rule evaluates.
<code>value</code>	Specifies the value against which the specified attribute is evaluated.

routing:less

Synopsis

```
<routing:less contextName="..."
              contextAttributeName="..."
              value="..." />
```

Description

The `routing:less` element is a child of a `routing:transportAttribute` element. It defines a rule that is triggered when the value of the specified attribute is less than the value given. It applies to numeric attributes.

Attributes

The `routing:less` element has the following attributes:

<code>contextName</code>	Specifies the QName of the context in which the desired transport attributes are stored.
<code>contextAttributeName</code>	Specifies the QName of the transport attribute the rule evaluates.
<code>value</code>	Specifies the value against which the specified attribute is evaluated.

routing:startswith

Synopsis

```
<routing:startswith contextName="..."
                    contextAttributeName="..."
                    value="..."
                    ignorecase="..." />
```

Description

The `routing:startswith` element is a child of a [routing:transportAttribute](#) element. It applies to string attributes and tests whether the attribute starts with the specified value.

Attributes

The `routing:startswith` element has the following attributes:

<code>contextName</code>	Specifies the QName of the context in which the desired transport attributes are stored.
<code>contextAttributeName</code>	Specifies the QName of the transport attribute the rule evaluates.
<code>value</code>	Specifies the value against which the specified attribute is evaluated.
<code>ignorecase</code>	Specifies whether the case of characters in a string are ignored. The default is no; case is considered when evaluating string data.

routing:endswith

Synopsis

```
<routing:endswith contextName="..."
                  contextAttributeName="..."
                  value="..."
                  ignorecase="..." />
```

Description

The `routing:endswith` element is a child of a [routing:transportAttribute](#) element. It applies to string attributes and tests whether the attribute ends with the specified value.

Attributes

The `routing:endswith` element has the following attributes:

<code>contextName</code>	Specifies the QName of the context in which the desired transport attributes are stored.
--------------------------	--

<code>contextAttributeName</code>	Specifies the QName of the transport attribute the rule evaluates.
<code>value</code>	Specifies the value against which the specified attribute is evaluated.
<code>ignorecase</code>	Specifies whether the case of characters in a string are ignored. The default is no; case is considered when evaluating string data.

routing:contains

Synopsis

```
<routing:contains contextName="..."
                  contextAttributeName="..."
                  value="..."
                  ignorecase="..." />
```

Description

The `routing:contains` element is a child of a [routing:transportAttribute](#) element. It applies to string or list attributes. For strings, it tests whether the attribute contains the value. For lists, it tests whether the value is a member of the list.

Attributes

The `routing:contains` element has the following attributes:

<code>contextName</code>	Specifies the QName of the context in which the desired transport attributes are stored.
<code>contextAttributeName</code>	Specifies the QName of the transport attribute the rule evaluates.
<code>value</code>	Specifies the value against which the specified attribute is evaluated.
<code>ignorecase</code>	Specifies whether the case of characters in a string are ignored. The default is no; case is considered when evaluating string data.

routing:empty

Synopsis

```
<routing:empty contextName="..."
               contextAttributeName="..." />
```

Description

The `routing:empty` element is a child of a [routing:transportAttribute](#) element. It applies to string or list attributes. For lists, it tests whether the list is empty. For strings, it tests for an empty string.

Attributes

The `routing:empty` element has the following attributes:

<code>contextName</code>	Specifies the QName of the context in which the desired transport attributes are stored.
<code>contextAttributeName</code>	Specifies the QName of the transport attribute the rule evaluates.

`routing:nonempty`

Synopsis

```
<routing:nonempty contextName="..."
                  contextAttributeName="..." />
```

Description

The `routing:nonempty` element is a child of a `routing:transportAttribute` element. It applies to string or list attributes. For lists, it passes if the list is not empty. For strings, it passes if the string is not empty.

Attributes

The `routing:nonempty` element has the following attributes:

<code>contextName</code>	Specifies the QName of the context in which the desired transport attributes are stored.
<code>contextAttributeName</code>	Specifies the QName of the transport attribute the rule evaluates.

Transport Attribute Context Names

The `contextName` attribute is specified using the QName of the context in which the attribute is defined. The contexts shipped with Artix are described in [Table 18](#).

Table 18: *Context QNames*

Context QName	Details
<code>http-conf:HTTPServerIncomingContexts</code>	Contains the attributes for HTTP messages being received by a server.
<code>corba:corba_input_attributes</code>	Contains the data stored in the CORBA principle

Table 18: *Context QName*

Context QName	Details
mq:MQConnectionAttributes	Contains the attributes used to connect to an MQ queue.
mq:MQIncomingMessageAttributes	Contains the attributes in the message header of an MQ message.
bus-security	Contains the attributes used by the security service to secure your services.

Security

Artix uses a special WSDL extension element to specify security policies for endpoints.

Runtime Compatibility

The security extensions are only compatible with C++ runtime.

Namespace

The elements Artix uses for specifying security policies are defined in the `http://schemas.iona.com/bus/security` namespace. When defining security policies in an Artix contract your contract's definition element must have the following entry:

```
xmlns:bus-security="http://schemas.iona.com/bus/security"
```

`bus-security:security`

Synopsis

```
<bus-security:security enableSecurity="..."
    is2AuthorizationActionRoleMapping="..."
    enableAuthorization="..."
    authenticationCacheSize="..."
    authenticationCacheTimeout = "..."
    securityType="..."
    securityLevel="..."
```

```
authorizationRealm="..."
defaultPassword="..." />
```

Description

The `bus-security:security` element is a child of a WSDL port element. It's attributes specify security policies for the endpoint.

Attributes

The `bus-security:security` element has the following attributes:

<code>enableSecurity</code>	Specifies if the service should load the ASP plug-in. Default is <code>false</code> .
<code>is2AuthorizationActionRoleMapping</code>	Specifies the URL of the action role mapping file the Artix security framework uses to authenticate requests for this endpoint.
<code>enableAuthorization</code>	Specifies if the endpoint should use the Artix security framework for authentication. Default is <code>false</code> .
<code>enableSSO</code>	Specifies if the service can use single-sign on (SSO). Default is <code>false</code> .
<code>authenticationCacheSize</code>	Specifies the maximum number of credentials stored in the authentication cache. A value of <code>-1</code> (the default) means unlimited size. A value of <code>0</code> disables the cache.
<code>authenticationCacheTimeout</code>	Specifies the time (in seconds) after which a credential is considered stale. A value of <code>-1</code> (the default) means an infinite time-out. A value of <code>0</code> disables the cache.
<code>securityLevel</code>	Specifies the level from which security credentials are picked up. The following options are supported by the Artix security framework: <ul style="list-style-type: none"> • <code>MESSAGE_LEVEL</code>—Get security information from the transport header. This is the default. • <code>REQUEST_LEVEL</code>—Get the security information from the message header.

<code>authorizationRealm</code>	Specifies the Artix authorization realm to which an Artix server belongs. The value of this variable determines which of a user's roles are considered when making an access control decision. The default is <code>IONAGlobalRealm</code> .
<code>defaultPassword</code>	Specifies the password to use on the server side when the client credentials originate either from a CORBA Principal (embedded in a SOAP header) or from a certificate subject. The default is <code>default_password</code> .

See also

For more information about Artix security policies, see [The Artix Security Guide](#).

Codeset Conversion

For transports that do not natively support codeset conversion Artix has the ability to perform codeset conversion.

Runtime Compatibility

The extension elements used to configure codeset conversion are only compatible with the C++ runtime.

Namespace

The elements Artix uses for defining codeset conversion rules are defined in the `http://schemas.iona.com/bus/i18n/context` namespace. When defining codeset conversion rules in an Artix contract your contract's definition element must have the following entry:

```
xmlns:i18n-context="http://schemas.iona.com/bus/i18n/context"
```

`i18n-context:client`

Synopsis

```
<i18n-context:client LocalCodeSet="..." OutboundCodeSet="..."  
    InboundCodeSet="..." />
```

Description

The `i18n-context:client` element is a child of a WSDL `port` element. It specifies codeset conversion rules for Artix endpoints that are acting as servers.

Attributes

The `i18n-context:client` element has the following attributes for defining how message codesets are converted:

<code>LocalCodeSet</code>	Specifies the client's native codeset. Default is the codeset specified by the local system's locale setting.
<code>OutboundCodeSet</code>	Specifies the codeset into which requests are converted. Default is the codeset specified in <code>LocalCodeSet</code> .
<code>InboundCodeSet</code>	Specifies the codeset into which replies are converted. Default is the codeset specified in <code>OutboundCodeSet</code> .

i18n-context:server**Synopsis**

```
<i18n-context:server LocalCodeSet="..." OutboundCodeSet="..."
    InboundCodeSet="..." />
```

Description

The `i18n-context:server` element is a child of a WSDL `port` element. It specifies codeset conversion rules for Artix endpoints that are acting as servers.

Attributes

The `i18n-context:server` element has the following attributes for defining how message codesets are converted:

<code>LocalCodeSet</code>	Specifies the server's native codeset. Default is the codeset specified by the local system's locale setting.
<code>OutboundCodeSet</code>	Specifies the codeset into which replies are converted. Default is the codeset specified in <code>InboundCodeSet</code> .
<code>InboundCodeSet</code>	Specifies the codeset into which requests are converted. Default is the codeset specified in <code>LocalCodeSet</code> .

Index

A

- adding a SOAP header 37, 45
- arrays
 - mapping to a fixed binding 92
 - mapping to a tagged binding 99
 - mapping to a TibrvMsg 111
 - mapping to CORBA 67
- Artix contexts
 - using in a TibrvMsg 117
- Artix reference
 - mapping to CORBA 73
- attribute based routing 200

B

- bus-security:security 208
 - authenticationCacheSize attribute 208
 - authenticationCacheTimeout attribute 208
 - authorizationRealm attribute 209
 - defaultPassword attribute 209
 - enableAuthorization attribute 208
 - enableSecurity attribute 208
 - enableSSO attribute 208
 - is2AuthorizationActionRoleMapping attribute 208
 - securityLevel attribute 208

C

- choice complexType
 - mapping to a fixed binding 90
 - mapping to a tagged binding 101
- complex types
 - mapping to a TibrvMsg 115
 - mapping to CORBA 60
- corba:address 145
 - location attribute 146
- corba:alias 66
 - name attribute 66
 - repositoryID attribute 66
 - type attribute 66
- corba:anonsequence 70
 - bound attribute 70
 - elemtype attribute 70
 - name attribute 70

- type attribute 70
- corba:array 67
 - bound attribute 67
 - elemtype attribute 67
 - name attribute 67
 - repositoryID attribute 67
 - type attribute 67
- corba:binding 56
 - bases attribute 56
 - repositoryID attribute 56
- corba:case 65
 - label attribute 65
- corba:enumerator 62
- corba:exception 69
 - name attribute 69
 - repositoryID attribute 69
 - type attribute 69
- corba:fixed 63
 - digits attribute 63
 - name attribute 63
 - repositoryID attribute 63
 - scale attribute 63
 - type attribute 63
- corba:member 60
 - idltype attribute 60
 - name attribute 60
- corba:object
 - binding attribute 74
 - name attribute 74
 - repositoryID attribute 74
 - type attribute 74
- corba:operation 56
 - name attribute 57
- corba:param 57
 - idltype attribute 57
 - mode attribute 57
 - name attribute 57
- corba:policy 146
 - persistent attribute 147
 - poaname attribute 147
 - serviceid attribute 147
- corba:raises 58
 - exception attribute 58

- corba:return 58
 - idtype attribute 58
 - name attribute 58
- corba:sequence 68
 - bound attribute 68
 - elemtype attribute 68
 - name attribute 68
 - repositoryID attribute 68
- corba:typeMapping 59
 - targetNamespace attribute 59
- corba:union 64
 - discriminator attribute 64
 - name attribute 64
 - repositoryID attribute 64
 - type attribute 64
- corba:unionbranch 64
 - default attribute 65
 - idtype attribute 64
 - name attribute 64

D

- defining a fixed message body 84
- defining a tagged message body 97
- defining a TibrvMsg 115
- durable subscriptions 179

E

- enumerations
 - mapping to a fixed binding 88
 - mapping to a tagged binding 98
 - mapping to CORBA 61
- exceptions
 - mapping to CORBA 58, 69
 - mapping to SOAP 38, 46

F

- failover routing 198
- fanout routing 198
- fixed:binding 83
 - encoding attribute 84
 - justification attribute 84
 - padHexCode attribute 84
- fixed:body 84
 - encoding attribute 85
 - justification attribute 85
 - padHexCode attribute 85
- fixed:case 90
 - fixedValue attribute 91

- name attribute 91
- fixed:choice 90
 - discriminatorName attribute 90
 - name attribute 90
- fixed:enumeration 88
 - fixedValue attribute 89
 - value attribute 89
- fixed:field 85
 - bindingOnly attribute 86
 - fixedValue attribute 86
 - format attribute 86
 - justification attribute 86
 - name attribute 85
 - size attribute 86
- fixed:operation 84
 - discriminator attribute 84
- fixed:sequence 92
 - counterName attribute 93
 - name attribute 93
 - occurs attribute 93
- ftp:port 192
 - connectMode 192
 - host 192
 - port 192
 - replyLocation 192
 - requestLocation 192
 - scanInsterval 192
- ftp:properties 192
- ftp:property 193
 - name 193
 - value 193

H

- http:address 126
 - location attribute 126
- http-conf:client 127
 - Accept attribute 133
 - AcceptEncoding attribute 135
 - AcceptLanguage attribute 134
 - Authorization attribute 133
 - AuthorizationType attribute 133
 - AutoRedirect attribute 128
 - BrowserType attribute 140
 - CacheControl attribute 137
 - cache-extension directive 138
 - max-age directive 137
 - max-stale directive 138
 - min-fresh directive 138
 - no-cache directive 137

- no-store directive 137
- no-transform directive 138
- only-if-cached directive 138
- ClientCertificate attribute 129
- ClientCertificateChain attribute 129
- ClientPrivateKey attribute 130
- ClientPrivateKeyPassword attribute 130
- ConnectionAttempts attribute 129
- Connection attribute 137
- ContentType attribute 128
- Cookie attribute 129
- Host attribute 136
- Password attribute 128
- ProxyAuthorization attribute 142
- ProxyAuthorizationType attribute 142
- ProxyPassword attribute 129
- ProxyServer attribute 142
- ProxyUserName attribute 129
- ReceiveTimeout attribute 128
- Referer attribute 140
- SendTimeout attribute 128
- TrustedRootCertificate attribute 130
- UserName attribute 128
- UseSecureSockets attribute 143
- http-conf:server 130
 - CacheControl attribute 137
 - cache-extension directive 140
 - max-age directive 139
 - must-revalidate directive 139
 - no-cache directive 139
 - no-store directive 139
 - no-transform directive 139
 - private directive 139
 - proxy-revalidate directive 139
 - public directive 139
 - s-maxage directive 140
 - ContentEncoding attribute 136
 - ContentLocation attribute 131
 - ContentType attribute 131
 - HonorKeepAlive attribute 131
 - ReceiveTimeout attribute 130
 - RedirectURL attribute 143
 - SendTimeout attribute 130
 - ServerCertificate 132
 - ServerCertificateChain 143
 - ServerPrivateKey attribute 132
 - ServerPrivateKeyPassword attribute 132
 - ServerType attribute 131
 - SuppressClientReceiveErrors attribute 131

- SuppressClientSendErrors attribute 131
- TrustedRootCertificate attribute 132
- UseSecureSockets attribute 143

I

- i18n-context:client 212
 - InboundCodeSet 212
 - LocalCodeSet 212
 - OutboundCodeSet 212
- i18n-context:server 212
 - InboundCodeSet 212
 - LocalCodeSet 212
 - OutboundCodeSet 212
- IDL types
 - fixed 62
 - Object 73
 - sequence 68
 - typedef 66
- iiop:address 149
 - location attribute 150
- iiop:payload 150
 - type attribute 150
- iiop:policy 151
 - persistent attribute 151
 - poaname attribute 151
 - serviceid attribute 151
- IOR 145, 149

J

- jms:address 176
 - connectionPassword attribute 177
 - connectionUserName attribute 177
 - destinationStyle attribute 176
 - jmsDestinationName attribute 177
 - jmsReplyDestinationName 177
 - jndiConnectionFactoryName attribute 176
 - jndiDestinationName attribute 176
 - jndiReplyDestinationName 177
- jms:client 178
 - messageType attribute 178
- jms:JMSNamingProperty 177
 - name attribute 177
 - value attribute 177
- jms:server 178
 - durableSubscriberName attribute 179
 - messageSelector attribute 179
 - transactional attribute 179
 - useMessageIDAsCorrelationID attribute 179

JNDI

- connection factory 176

L

- load balancing 198

M

- message broadcasting 198

- mime:content 50

- part attribute 51
 - type attribute 51

- mime:multipartRelated 50

- mime:part 50

- name attribute 50

- mq:client 155

- AccessMode attribute 165

- AccountingToken attribute 156

- AliasQueueName attribute 161

- ApplicationData attribute 156

- ApplicationIdData attribute 156

- ApplicationOriginData attribute 156

- ConnectionFastPath attribute 155

- ConnectionName attribute 155

- ConnectionReusable attribute 155

- CorrelationId attribute 156

- CorrelationStyle attribute 164

- Delivery attribute 166

- Format attribute 170

- MessageExpiry attribute 156

- MessageId attribute 156

- MessagePriority attribute 166

- ModelQueueName attribute 156

- QueueManager attribute 155

- QueueName attribute 155

- ReplyQueueManager attribute 155

- ReplyQueueName attribute 155

- ReportOption attribute 168

- Server_Client attribute 160

- Timeout attribute 156

- Transactional attribute 167

- UsageStyle attribute 163

- UserIdentification attribute 156

- mq:server 157

- AccessMode attribute 165

- AccountingToken attribute 159

- ApplicationData attribute 158

- ApplicationOriginData attribute 159

- ConnectionFastPath attribute 158

- ConnectionName attribute 157

- ConnectionReusable attribute 158

- CorrelationId attribute 158

- CorrelationStyle attribute 164

- Delivery attribute 166

- Format attribute 170

- MessageExpiry attribute 158

- MessageId attribute 158

- MessagePriority attribute 166

- ModelQueueName attribute 157

- PropagateTransactions attributes 159

- QueueManager attribute 157

- QueueName attribute 157

- ReplyQueueManager attribute 157

- ReplyQueueName attribute 157

- ReportOption attribute 168

- Server_Client attribute 160

- Timeout attribute 158

- Transactional attribute 167

- UsageStyle attribute 163

P

- passthru:binding 122

- POA policies 146, 151

- port address

- HTTP 126

- primitive types

- mapping to a fixed binding 85

- mapping to a tagged binding 98

- mapping to a TibrvMsg 106, 116

- mapping to CORBA 54

- mapping to FML 80

R

- reply queue

- queue manager 155, 157

- queue name 155, 157

- request queue

- queue manager 155, 157

- queue name 155, 157

- routing:contains 204

- contextAttributeName attribute 204

- contextName attribute 204

- ignorecase attribute 204

- value attribute 204

- routing:destination 199

- port attribute 200

- route attribute 200

- service attribute 200
- value attribute 200
- routing:empty 204
 - contextAttributeName attribute 205
 - contextName attribute 205
- routing:endswith 203
 - contextAttributeName attribute 204
 - contextName attribute 203
 - ignorecase attribute 204
 - value attribute 204
- routing:equals 201
 - contextAttributeName attribute 201
 - contextName attribute 201
 - ignorecase attribute 201
 - value attribute 201
- routing:expression 197
 - evaluator attribute 198
 - name attribute 198
- routing:greater 202
 - contextAttributeName attribute 202
 - contextName attribute 202
 - value attribute 202
- routing:less 202
 - contextAttributeName attribute 202
 - contextName attribute 202
 - value attribute 202
- routing:nonempty 205
 - contextAttributeName attribute 205
 - contextName attribute 205
- routing:query 199
- routing:route 198
 - multiRoute attribute 198
 - failover 198
 - fanout 198
 - loadBalance 198
 - name attribute 198
- routing:source 199
 - port attribute 199
 - service attribute 199
- routing:startswith 203
 - contextAttributeName attribute 203
 - contextName attribute 203
 - ignorecase attribute 203
 - value attribute 203
- routing:transportAttribute 200

S

- sequence complexType
 - mapping to a fixed binding 92

- mapping to a tagged binding 99
- service failover 198
- soap:address 126
 - location attribute 126
- soap:binding 33
 - style attribute 33
 - transport attribute 34
- soap:body 35
 - encodingStyle attribute 37
 - namespace attribute 37
 - parts attribute 37
 - use attribute 36
 - encoded 36
 - literal 36
- soap:fault 38
 - name attribute 39
 - use attribute 39
 - encoded 36
 - literal 36
- soap:header 37
 - encodingStyle attribute 38
 - message attribute 38
 - namespace attribute 38
 - part attribute 38
 - use attribute 38, 46
 - encoded 36
 - literal 36
- soap:operation 34
 - soapAction attribute 35
 - style attribute 35
- specifying a password
 - HTTP 128
- specifying a user name
 - HTTP 128

T

- tagged:binding 95, 96, 122
 - fieldNameValueSeparator attribute 96
 - fieldSeparator attribute 96
 - flattened attribute 96
 - ignoreCase attribute 97
 - ignoreUnknownElements attribute 97
 - messageEnd attribute 96
 - messageStart attribute 96
 - scopeType attribute 96
 - selfDescribing attribute 96
 - unscopedArrayElement attribute 96
- tagged:body 97
- tagged:case 102

- name attribute 102
- tagged:choice 101
 - alias attribute 101
 - discriminatorName attribute 101
 - name attribute 101
- tagged:enumeration 98
 - value attribute 98
- tagged:field 98
 - alias attribute 98
 - name attribute 98
- tagged:operation 97
 - discriminator attribute 97
 - discriminatorStyle attribute 97
- tagged:sequence 99
 - alias attribute 100
 - name attribute 100
 - occurs attribute 100
- tibrv:array 111
 - elementName attribute 112
 - integralAsSingleField attribute 112
 - loadSize attribute 112
 - sizeName attribute 112
- tibrv:binding 107
 - stringAsOpaque attribute 108
 - stringEncoding attribute 108
- tibrv:context 117
- tibrv:field 116
 - alias attribute 116
 - element attribute 116
 - id attribute 116
 - maxOccurs attribute 116
 - minOccurs attribute 116
 - name attribute 116
 - type attribute 116
 - value attribute 116
- tibrv:input 109
 - messageNameFieldPath attribute 109
 - messageNameFieldValue attribute 109
 - stringAsOpaque attribute 110
 - stringEncoding attribute 109
- tibrv:msg 115
 - alias attribute 115
 - element attribute 115
 - id attribute 115
 - maxOccurs attribute 115
 - minOccurs attribute 115
 - name attribute 115
- tibrv:operation 108
- tibrv:output 110
 - messageNameFieldPath attribute 111
 - messageNameFieldValue attribute 111
 - stringAsOpaque attribute 111
 - stringEncoding attribute 111
- tibrv:port 183
 - bindingType attribute 187
 - callbackLevel attribute 187
 - clientSubject attribute 183
 - cmListenerCancelAgreements attribute 185
 - cmQueueTransportCompleteTime attribute 186
 - cmQueueTransportSchedulerActivation attribute 186
 - cmQueueTransportSchedulerHeartbeat attribute 185
 - cmQueueTransportSchedulerWeight attribute 185
 - cmQueueTransportServerName attribute 189
 - cmQueueTransportWorkerTasks attribute 185
 - cmQueueTransportWorkerWeight attribute 185
 - cmSupport attribute 184
 - cmTransportClientName attribute 184
 - cmTransportDefaultTimeLimit attribute 185
 - cmTransportLedgerName attribute 184
 - cmTransportRelayAgent attribute 185
 - cmTransportRequestOld attribute 184
 - cmTransportServerName attribute 188
 - cmTransportSyncLedger attribute 185
 - responseDispatchTimeout attribute 188
 - serverSubject attribute 183
 - transportBatchMode attribute 184
 - transportDaemon attribute 184
 - transportNetwork attribute 188
 - transportService attribute 188
- timeouts
 - HTTP 128
 - MQ 156, 158
- transactions
 - MQ 167
- tuxedo:binding 80
- tuxedo:field 81
 - id attribute 81
 - name attribute 81
- tuxedo:fieldTable 81
 - type attribute 81
- tuxedo:input 174
 - operation attribute 174
- tuxedo:operation 81
- tuxedo:server 174
- tuxedo:service 174
 - name attribute 174

U

- unions
 - mapping to a fixed binding 90
 - mapping to a tagged binding 101
 - mapping to CORBA 64

W

- wsoap12:address 126
 - location attribute 126
- wsoap12:binding 41
 - style attribute 41
 - transport attribute 42
- wsoap12:body 43
 - encodingStyle attribute 45
 - namespace attribute 45
 - parts attribute 45
 - use attribute 44
 - literal 44
- wsoap12:fault 46
 - name attribute 47
 - namespace attribute 47
 - use attribute 47

- literal 44
- wsoap12:header 45
 - encodingStyle attribute 46
 - message attribute 45
 - namespace attribute 46
 - part attribute 45
 - use attribute
 - literal 44
- wsoap12:operation 43
 - soapAction attribute 43
 - soapActionRequired attribute 43
 - style attribute 43
- wsoap12/
 - fault
 - encodingStyle attribute 47

X

- xformat:binding 120
 - rootNode attribute 120
- xformat:body 120
 - rootNode attribute 120

