

# Network Requirements for ISIS and Interplay Production

David Shephard EURING CEng MIET CCDP® CCNP® CCIP®  
Senior Network Solutions Architect  
01 October, 2012



This document is available from:  
Network Requirements for ISIS 7000 and Interplay.  
[http://avid.force.com/pkb/articles/en\\_US/Compatibility/en244197](http://avid.force.com/pkb/articles/en_US/Compatibility/en244197)  
*Intended audience: General distribution*

---

## Abstract

This document outlines the fundamental requirements for ISIS 7000 solutions with Interplay Production. It is intended to provide a summary of many documents. And set out minimum requirements where such direction is not explicitly documented, but experience from existing installations is applicable. The document content may be updated in line with product S/W and H/W releases or when other content is added not in direct relation to a recent software releases. Externally available URLs will be provided where possible.

*Note: All references to ISIS 7000 equates to the new name of ISIS7000. Absolute references to ISIS5000 apply to ISIS 5000 only*

This document can be shared with customers and used for SoW content.

© Avid Technology (Europe) Ltd. This document is the property of Avid. The information contained in this document has been provided to the intended recipients for evaluation purposes only. The information contained in this document should not be discussed with any other party or persons without the express prior written permission of Avid. If the intended recipient does not accept these terms, this document and any copies should be returned to the nearest Avid office. If you are not the intended recipient, employee or agent you are hereby notified that any dissemination or copying of this document is strictly prohibited. If you have received this document in error, please return it to the nearest Avid Technology office ([www.avid.com](http://www.avid.com)).



# Table of Contents

<b>ABSTRACT .....</b>	<b>1</b>
<i>Recent Revision history.....</i>	<i>6</i>
<b>1.0 ISIS REQUIREMENTS.....</b>	<b>8</b>
1.0.1 <i>Quality of Service – Latency and Jitter .....</i>	<i>8</i>
1.0.2 <i>Latency impact on ISIS network traffic .....</i>	<i>9</i>
1.0.3 <i>Blueprint Architecture for Cisco and ISIS .....</i>	<i>9</i>
1.1 <b>THE BORDER SWITCH .....</b>	<b>9</b>
1.2 <b>ZONE DESCRIPTIONS.....</b>	<b>10</b>
1.3 <b>QUALIFIED SWITCHES .....</b>	<b>11</b>
1.3.1 <i>Using Fast Ethernet .....</i>	<i>12</i>
1.3.2 <i>Using SMC 8724ML3.....</i>	<i>13</i>
1.4 <b>APPROVED SWITCHES .....</b>	<b>13</b>
1.4.1 <i>Cisco 6500.....</i>	<i>13</i>
1.4.2 <i>Use of WS-X6708-10G-3C.....</i>	<i>14</i>
1.4.3 <i>Use of WS-X6716-10G-3C.....</i>	<i>15</i>
1.4.4 <i>SUP 2T and 6800 /6900 series I/O module.....</i>	<i>15</i>
1.4.5 <i>Cisco Catalyst 4500 Classic.....</i>	<i>15</i>
1.4.6 <i>Cisco Catalyst 4500 Enhanced .....</i>	<i>16</i>
1.4.7 <i>Cisco Nexus 7000.....</i>	<i>17</i>
1.4.9 <i>Foundry/Brocade FESX 448/648 .....</i>	<i>18</i>
1.4.10 <i>Foundry/Brocade SuperX.....</i>	<i>19</i>
1.4.11 <i>Foundry/Brocade Big Iron RX .....</i>	<i>20</i>
1.4.12 <i>Arista Networks 7048.....</i>	<i>21</i>
1.4.13 <i>Cisco Catalyst C4500-X.....</i>	<i>21</i>
1.5 <b>NON APPROVED SWITCHES, MODULES AND VOIP.....</b>	<b>22</b>
1.5.1 <i>Cisco Catalyst 3750 .....</i>	<i>22</i>
1.5.2 <i>Cisco Nexus 5500/5000/2000.....</i>	<i>22</i>
1.5.4 <i>4908 10G module for Cisco Catalyst 4900M.....</i>	<i>23</i>
1.5.5 <i>Juniper EX3200 and EX4200.....</i>	<i>23</i>
1.5.6 <i>Brocade/Foundry NetIron MLX .....</i>	<i>24</i>
1.5.7 <i>Switch Buffering architectures and limitations .....</i>	<i>25</i>
1.5.8 <i>Inline VoIP device.....</i>	<i>26</i>
1.6 <b>NETWORK INTERFACE CARD REQUIREMENTS .....</b>	<b>26</b>
1.6.1 <i>Using Fast Ethernet .....</i>	<i>27</i>
1.6.2 <i>When not to use the Intel PRO/1000M or Intel Pro/1000P NIC .....</i>	<i>28</i>
1.6.3 <i>HP xw8600 &amp; Z800 &amp; Z400 Workstations and Broadcom Ethernet Connections.....</i>	<i>29</i>
1.6.4 <i>Setting descriptors.....</i>	<i>29</i>
1.6.5 <i>10G network interfaces for Ultra High Resolution Clients.....</i>	<i>30</i>
1.6.7 <i>LAN on Motherboard.....</i>	<i>31</i>
1.6.8 <i>Intel Pro 1000 CT gigabit adapter.....</i>	<i>31</i>
1.6.9 <i>Avid Slot Configuration guide.....</i>	<i>31</i>
1.7 <b>DNS .....</b>	<b>31</b>
1.7.1 <i>DNS naming conventions .....</i>	<i>32</i>
1.8 <b>CABLE REQUIREMENTS.....</b>	<b>32</b>
1.8.1 <i>Copper cabling for Gigabit Ethernet .....</i>	<i>34</i>
1.8.2 <i>Fibre Optic cabling for 10 Gigabit Ethernet.....</i>	<i>35</i>
<i>Corning® InfiniCor® multimode fibers.....</i>	<i>37</i>
<i>SMF-28e® fiber .....</i>	<i>37</i>
1.8.3 <i>Fibre Optic Transceivers for 10 Gigabit Ethernet.....</i>	<i>37</i>
1.8.4 <i>Media Converters for Gigabit Ethernet .....</i>	<i>38</i>

1.8.5 Patching for Copper Structured Cabling .....	38
1.9 IP REQUIREMENTS .....	40
1.9.1 Ranges required .....	40
1.9.2 Default IP Ranges .....	41
1.9.3 VLAN numbering.....	43
1.9.4 Routed Interconnecting Networks .....	43
1.9.5 Interconnecting Networks example .....	43
1.9.6 Using Static Routes and HSRP.....	44
1.9.7 Routing protocols.....	45
1.9.8 ISIS 5000 IP address use.....	46
1.10 MAN/WAN CONNECTIONS .....	46
1.10.1 MAN - Example deployment.....	48
1.10.2 MAN - Proven deployment .....	48
1.11 DHCP.....	48
1.12 10G LINK AGGREGATION .....	49
1.13 DEPLOYING TRANSFER MANAGER.....	50
1.14 JUMBO FRAMES AND LEGACY APPLICATIONS .....	50
1.15 AVID LOW RES ENCODER .....	51
1.16 RESILIENT FIRST HOP PROTOCOL .....	51
1.17 INTERSWITCH LINK FOR RESILIENT CONFIGURATIONS.....	52
1.18 RSTP SETTINGS FOR FHRP IMPLEMENTATIONS.....	52
<b>2.0 INTERPLAY PRODUCTION REQUIREMENTS.....</b>	<b>54</b>
2.1 TO MULTICAST OR NOT TO MULTICAST .....	54
2.1.1 Multicast repeater - LEGACY.....	55
2.1.2 Direct Client Configuration .....	55
2.1.3 ALL Client Configuration Unicast.....	56
2.2 MULTICAST COMMANDS .....	56
2.2.1 Cisco Commands for Multicast .....	56
2.2.1 Foundry Cisco Commands for Multicast .....	57
2.2.3 ALL Client Configuration Unicast.....	58
2.3 DNS .....	58
2.3.1 Why is FQDN resolution required? .....	59
2.4 ACTIVE DIRECTORY WITH INTERPLAY PRODUCTION CLUSTER.....	59
2.4.1 Avid Interplay Production Active Directory Considerations .....	60
2.5 TIME-CODE AND NTP .....	62
2.5.1 Time Synchronisation for Avid Interplay™ systems .....	62
2.6 INTERPLAY PRODUCTION ASSIST BROWSE RESOLUTION.....	63
2.7 DHCP.....	63
2.8 STREAMING SERVER DEPLOYMENT PRACTICES .....	63
2.8.1 Network Zones and DNS.....	63
2.8.2 Network Requirements for Interplay Access streaming clients.....	64
2.8.3 Firewall Parameters for Interplay Stream Server Clients .....	64
2.8.4 Firewall Parameters for Interplay Streaming Server Clients .....	65
2.8.5 Supported Config .....	65
2.9 INTERPLAY PRODUCTION COPY SERVER.....	65
2.10 INTERPLAY PRODUCTION MOVE SERVER.....	68
<b>3.0 ENHANCING NETWORK PERFORMANCE.....</b>	<b>71</b>
3.1 TCP WINDOW SIZING .....	71
3.2 USEFUL KNOWLEDGE BASE ARTICLES .....	72
3.2.1 Starbucks Fix for ISIS v1.0-1.4 .....	72
<b>4.0 PC AND MAC REQUIREMENTS .....</b>	<b>72</b>
4.0.1 SEPTEMBER 2012 URLs.....	72
4.0.2 AUGUST 2011 URLs .....	72
4.1 CUSTOMER PROVIDED PLATFORMS .....	73
4.2 CUSTOMER TESTED PLATFORMS - 2007.....	73
4.3 IMAGING PC CLIENTS .....	74
4.4 ULTRA HIGH RESOLUTION CLIENTS .....	74



<b>5.0 NETWORK DESIGNS .....</b>	<b>75</b>
5.0.1 Cisco 6500.....	76
5.0.2 Cisco 6500 and 4500.....	77
5.0.3 Cisco 4948 with cascaded 3750.....	78
5.0.4 Foundry RX-8 core with FESX Edge .....	79
5.0.6 Zone 3 Mezzanine Network conceptual diagram .....	81
5.0.7 4900M example#1 – Reference Architecture .....	82
5.0.8 4900M example#2 – Reference Architecture .....	83
5.0.9 Nexus 7000 core & C4948E edge .....	83
5.0.10 Nexus 7000 core & C4948E edge – Dual stack ISIS.....	84
5.1 BUFFERING .....	86
5.2 CONNECTION VIA IP PHONES – NOT RECOMMENDED .....	86
5.3 USING A DUAL NETWORK CONNECTION .....	86
5.4 USING A TEAMED NETWORK CONNECTION .....	87
<b>6.0 FIREWALL ISIS AND INTERPLAY PRODUCTION.....</b>	<b>88</b>
6.1 FIRST UNDERSTAND ISIS TRAFFIC .....	89
6.2 NEXT UNDERSTAND LATENCY .....	89
6.3 FIREWALL PROCESS .....	90
6.4 WHAT ABOUT INTERPLAY PRODUCTION? .....	91
6.5 WHAT PORTS ARE USED?.....	91
6.6 WHY DID ISIS PORTS USED CHANGE IN ISIS V1.4? .....	91
6.7 SUCCESSFULLY TESTED FIREWALLS.....	91
6.7.1 Juniper SRX 3400.....	92
6.7.2 Cisco ASA 5500-40 .....	92
6.7.3 Cisco FWSM for Catalyst 6500 – Limited suitability.....	92
<b>7.0 SECURITY RECOMMENDATIONS.....</b>	<b>92</b>
7.1 APPLYING SECURITY IN NETWORK DESIGN?.....	93
7.1.1 Mezzanine network.....	93
7.1.2 Using VMWARE.....	95
7.2 INTERNET CONNECTIVITY RESTRICTIONS? .....	97
<b>8.0 NETWORK MANAGEMENT AND MONITORING.....</b>	<b>97</b>
<b>9.0 DNXHD IN ZONE 3 AND 4 .....</b>	<b>102</b>
9.1 TEST SETUP: .....	103
9.1.1 Test Equipment.....	103
9.1.2 Test Results .....	103
9.2 TEST SUMMARY.....	104
<b>APPENDIX A. HOW TO INTEGRATE INTERPLAY PRODUCTION ENGINE IN TRUSTED DOMAIN ENVIRONMENTS .....</b>	<b>105</b>
SCOPE.....	105
SCENARIO .....	105
CONFIGURATION REQUIREMENTS .....	105
CONFIGURATION STEPS.....	106
A. Standalone Interplay Production Engine.....	106
B. Clustered Interplay Production Engine .....	106
<b>APPENDIX B. SWITCH CONFIGURATION TIPS &amp; GOOD PRACTICES.....</b>	<b>107</b>
B.1 DOCUMENT YOUR CONFIGS WITH DESCRIPTIONS .....	107
B.2 SETTING SPANNING TREE TO RAPID SPANNING TREE.....	107
B2.1 Spanning tree cost.....	108
B.2.2 Spanning Cost type .....	108
B.3 SET PRIMARY SWITCH AS STP MASTER ROOT PRIMARY .....	109
B.4. SET SECONDARY SWITCH AS STP ROOT SECONDARY .....	110
B.5 DEPLOY BPDU GUARD ON ALL PORTS THAT USE PORTFAST.....	110
B5.1 Use ROOT GUARD on any interfaces that cascade to other switches .....	111
B.6 USE THE NO SHUTDOWN COMMAND ON ALL VLANS .....	112

B.7 USE THE SHUTDOWN COMMAND ON ALL UNUSED INTERFACES.....	112
B.8 ENABLE SECRET .....	112
B.9 PASSWORD ENCRYPTION .....	112
B.10 ENABLE TELNET .....	113
B.11 ENABLE SYNCHRONOUS LOGGING .....	113
B.12 GET PUTTY 0.06.....	113
B.13 LOGGING.....	113
B.14 USING A SYSLOG SERVER .....	114
B.15 TIMESTAMPS .....	114
B.16 SETTING THE TIME .....	115
B.17 SHOW TECH SUPPORT .....	115
<i>B16.2 What is listed?</i> .....	116
<i>B17.2 Show tech-support - CAVEATS</i> .....	116
<i>B17.3 How long does it take?</i> .....	116
B.18 HANDOVER PRACTICES .....	117
B.19 CISCO CATALYST 49XX SETTING OF THE CONFIG REGISTER .....	117
<b>APPENDIX C INTERPLAY CENTRAL AND KEMP LOAD BALANCER.....</b>	<b>118</b>
<b>APPENDIX D FAULT FINDING TIPS – TO BE ADDED .....</b>	<b>119</b>
<b>APPENDIX E FULL REVISION HISTORY .....</b>	<b>119</b>
<i>Revision history</i> .....	119

## Table of Figures

Figure 1 default descriptor allocations of 512TX and 256RX.....	30
Figure 2 descriptor allocations of 1024TX and 1024RX.....	30
Figure 3 - Avid Multicast Repeater for Interplay Production.....	55
Figure 4 - Copy Server in Zone 1 - Preferred .....	66
Figure 5 - Copy Server in Zone 2 .....	67
Figure 6 - Copy server in Zone 1 and Zone 2 .....	68
Figure 7 - MOVE server in Zone 1 .....	69
Figure 8 - Move Server in Zone 2.....	70
Figure 9 - Cisco 6500 Example .....	76
Figure 10 - Cisco 6500 and 4500 example .....	77
Figure 11 - Cascaded 3750G Example .....	78
Figure 12 – Foundry Example RX & FESX.....	79
Figure 13 - Foundry Super X and RX MRP Core .....	80
Figure 14 - High level plan of Mezzanine network structure .....	81
Figure 15 - Mezzanine network structure with Zone 3.1 and aggregated links.....	82
Figure 16 - Extended Mezzanine network structure with Zone 3.1 and aggregated links.....	83
Figure 17 - Nexus 7000 core & C4948E edge.....	84
Figure 18 - Nexus 7000 core & C4948E edge Dual stack ISIS.....	85
Figure 19 - AFT and SFT teaming examples.....	88
Figure 20 - Zone 3 Mezzanine Network Example.....	94
Figure 21 - DNxHD in Zone 3 and 4 .....	102

<b>Additional Contributors:</b>		
	Ralf Puchner	Emmanuel Derosier
Neil Tindal	WG5-Demons	Flock Demons
Joe Vandenberg	Jason Sturgill	Jamie White

Network Requirements for ISIS 7000 and Interplay Production. This document is available from:

[http://avid.force.com/pkb/articles/en\\_US/Compatibility/en244197](http://avid.force.com/pkb/articles/en_US/Compatibility/en244197)

## Recent Revision history

Note Version for this document number DOES NOT directly correlate to ISIS or Interplay Production version

For Full Revision History see Appendix C at end of this document

<b>Version</b>	<b>Name</b>	<b>Date</b>	<b>Comment</b>
Initial Issue V1.0	David Shephard	04 July 2007	
1.9		18 AUG 2011	<p>Update sections 1.2 1.5.1, 1.3, 1.4.4,1.4.8, 1.4.9, 1.5.7, 1.6, 1.8, 1.92, 2.0</p> <p>Add section 1.8.5 Patch Panels</p> <p>Add section 5.4 Using a teamed network connection</p> <p>Update top tips Appendix B add B.2.1 STP costs long/short</p> <p>Section 1.4.7 ( previously 1.5.2) Cisco Nexus 7000 approval , Other 1.4.x section incremented</p> <p>ADD 1.4.10.1 SUPER X 10G QD settings</p> <p>ADD 1.5.6.1 Foundry/Brocade MLXe</p> <p>Update 2.8.3, 2.8.4 Interplay Stream/Streaming server.</p> <p>Added section 4.4 UHRC clients</p> <p>Added section 1.6.5, 1.66, 1.6.7</p>
1.10			<p>Updates to appendix C</p> <p>Add new section 1.0.2 Latency impact on ISIS network traffic.</p> <p>Update 1.5.2 Cisco Nexus 5500/5000/2000</p> <p>Add 1.6.9 Avid Configuration Guidelines</p> <p>Add HSRP static routing New section 1.9.6</p> <p>Add 1.18 RSTP settings for FHRP implementations.</p> <p>Update 1.9.7 Routing protocols.</p> <p>Add 1.4.12 Arista Networks 7048 – Approved switch</p> <p>Add 1.4.13 Cisco Catalyst C4500-X - Approved switch</p> <p>Add Appendix C how to configure routing for KEMP load balancer and DMS with Interplay Central Playback services</p> <p>Amend diagrams in Section 5.0.7/8</p>

Version	Name	Date	Comment
			<b>Update section 1.4.2 Use of WS-X6708-10G-3C</b> <b>Update Avid website references with new KB URLs (where available)</b> <b>Add Section 5.0.9 and 5.0.10 with Nexus 7000 examples</b>

## 1.0 ISIS Requirements

ISIS 7000 is a high function real time editing systems and placed extensive demand on a network infrastructure. The design of the solution is key, to the successful operation and user acceptance of any new broadcast deployment.

### 1.0.1 Quality of Service – Latency and Jitter

The real time nature of editing high bandwidth video in a collaborative environment means that tolerance for delay and jitter is small. The table below shows that 5mS is the maximum latency which should be considered acceptable.

Value	Behavior	Comments
0ms	System performs on test network as if locally attached	
5ms	Noticeable degradation in scrubbing performance, slight delay in play function (minimal)	RECOMMENDED Maximum Jitter and Latency - combined
10ms	Particularly noticeable delay in scrubbing, 1s delay from pressing play to material playing, may not be suitable for editors	USEABLE
20ms	More noticeable delay in scrubbing, 2.5s delay from pressing play to material playing – this would most likely be unsuitable for editors	UNSUITABLE
50ms	Unusable delay from pressing play, buffer ran out after 4-5 seconds and then started dropping frames	NOT USEABLE
100ms	system will not mount ISIS workspaces, reports network errors	NON FUNCTIONAL

Based on the tests performed to determine maximum fibre optic distances, 5ms is an acceptable latency; this translates to a distance of a connection of approx. 1000-1500km\* where it would be acceptable to the operator.

\*Given that the speed of light constant in a vacuum, 'c' is exactly 299,792,458 meters per second, the figure of 1 millisecond per 300km might be an accurate estimate for the purpose of latency calculation over distance

However, propagation speed in media is significantly lower than c, for glass roughly 1/2 - 2/3 of light speed in vacuum, depending on the refraction index of the media, so a figure of 1 millisecond per 200km is more appropriate .

Hence a round trip time (RTT) of 1 ms per 100KM is a working figure is applied to longer distances but this does not consider delays encountered by network equipment such optical/electrical translation and networks switches.

Jitter or the variation in latency is also a factor, but tends to have less of an impact than latency, 5mS of jitter added to 5mS of latency = 10mS of latency, and the performance of the client will suffer. However, the usability of the application is dependant upon the nature of the application, for example an Interplay Production Browse client being used to review material will be affected much less by latency than a NewsCutter or Media Composer client actively editing.



### 1.0.2 Latency impact on ISIS network traffic

Generally the effect of latency does not impact UDP traffic, however if there are upper layer transactional messages on-going between the end points, as is the case with ISIS protocols then latency will have an impact on those which in turn will have an impact of throughput, and hence will impact video editing operations.

To emulate the effects of write consolidation & read throughput via a link with different latency characteristics, Using PATH DIAG UNLIMITED 4MB I/O transfer was used to identify the maximum speed attained with varying degrees of latency as shown below.

A 1 minute pass was made and the average rate was recorded. The read file size was 1GB.

Latency	WRITE	%age hit	READ	%age hit	READ DURATION inc. writing file
Natural 0.5ms	71 MB/S	n/a	97MB/S	n/a	1:14
+1ms =1.5	67 MB/S	6%	81MB/S	16%	1:15 (+0:01)
+2ms =2.5	64 MB/S	9%	68MB/S	30%	1:16 (+0:02)
+4ms =4.5	58 MB/S	18%	52 MB/S	46%	1:17 (+0:03)
+8ms =8.5	48 MB/S	32%	35 MB/S	64%	1:22 (+0:08)
+16ms =16.5	33 MB/S	53%	21 MB/S	78%	1:28 (+0:16)
+32ms =32.5	21 MB/S	70%	12 MB/S	88%	1:46 (+0:32)

TESTING MARCH 2012: ISIS 2.4 server with ISIS 3.5 client.

This indicates that a write consolidate process on a circuit with approx. 5ms of latency will be 20% slower than normal, assuming no packet drop.

FPING was used to measure latency as this is accurate to 0.1ms and the extra granularity is required because the default window ping is only accurate to 1mS.

<http://www.kwakkelflap.com/fping.html>

### 1.0.3 Blueprint Architecture for Cisco and ISIS

Cisco and Avid have jointly published (April 2010) a document that provides recommendations to enable Broadcasters to deploy Enterprise Network Architectures supporting the Media Workflow Platform (MWP)

Best Practices for ISIS Networking in a Cisco environment is available from the Cisco Website and the Avid Website and the Avid Knowledge base URL:

[http://avid.force.com/pkb/articles/en\\_US/White\\_Paper/en362611](http://avid.force.com/pkb/articles/en_US/White_Paper/en362611)

### 1.1 The Border Switch

Most ISIS 7000 implementations require some communications between the ISIS VLANs and with the corporate network. A suitable layer 3 switch is required to provide this service. There are different designs which can be applied; however the list of tested devices and configurations is limited.





Avid uses the concept of Zones to define required levels of functionality and Quality of Service that must be provided. Zones are described in Section 1.2

## 1.2 Zone Descriptions

### ISIS Client vs. Zone Description

#### Zone-1 Client

- Connected to ISIS VLAN(s) via ISS 1Gb Port (direct connect)

#### Zone-2 Client

- Connected to ISIS VLAN(s) via 1Gb port on Avid qualified L2 Switch (non routed)

#### Zone-3 Client

- Connected to an Avid Qualified Layer-3 Switch (Router) with known QoS
- Traffic routed to ISIS (1 hop) and load balanced across ISIS VLANs (~60/40 ratio)

#### Zone-4 Client

- Connected to Customer's house network using Customers Edge/Core Switch with unknown QoS
- Traffic routed to ISIS (? Hops) and load balanced across ISIS VLANs (~60/40 ratio)

Support for different client types vary by Zone – for example in ISIS 7000 V1.x typical Zonal definitions might have deployed as:

- Zone-1: AirSpeed Payout, Transfer Manager
- Zone-2: AirSpeed Ingest, Editors,
- Zone-3: Interplay Production Engines, Instinct, Assist, Certain Editors (e.g. NC)  
Typically: DV25, DV50/IMX-50, MPEG-2 Proxy (2 Mb/s)
- Zone-4: Instinct, Assist  
Typically: DV25, MPEG-2 Proxy (2 Mb/s)

With ISIS V2.x hardware the fabric bandwidth has increase significantly and many devices can be used in Zone or Zone 3 depending on the application, and regardless of BW demand up to the limits of the ISS. The necessary planning is performed by a Network Consultant or for less complex scenarios a Project Engineer

Sometimes a smaller switch will be cascaded (down-linked from) off a larger switch, but still be within the QoS administration of Avid. One example of this would be a Cisco Catalyst 4900M connecting to ISIS with an aggregated 10G connection, with the Video client connecting via local Cisco Catalyst 4948 which connect back with 10G to the Cisco Catalyst 4900M. This could provide both Zone 2.1 and/or Zone 3.1 connections, and has been deployed successfully.

Zone 2 - a Gigabit Ethernet (L2) switch port with a direct 10G connection to ISIS

Zone 2.1 - a Gigabit Ethernet (L2) switch port with an indirect 1 hop 10G connection to ISIS

Zone 2.2 - a Gigabit Ethernet (L2) switch port with an indirect 2 hop 10G connection to ISIS

Zone 3 - a Gigabit Ethernet (L3) switch port with a direct 10G connection to ISIS

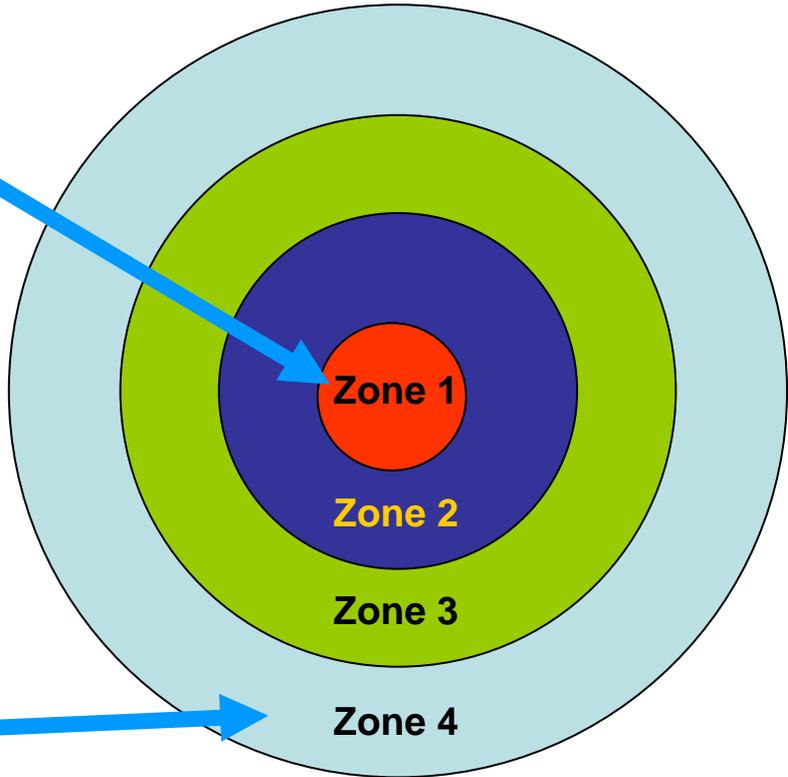
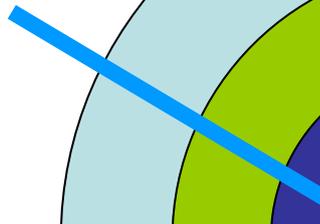
Zone 3.1 - a Gigabit Ethernet (L3) switch port with an indirect 1 hop 10G connection to ISIS

Zone 3.2 - a Gigabit Ethernet (L3) switch port with an indirect 2 hop 10G connection to ISIS

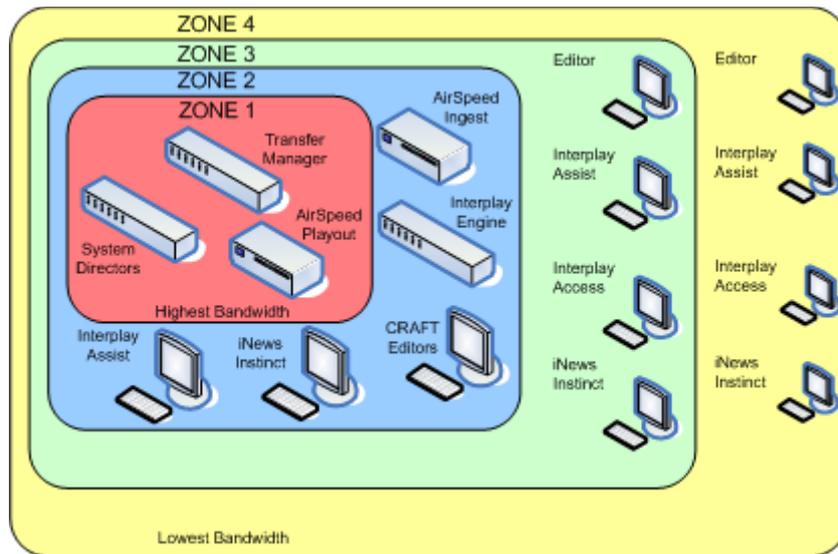
Note Zone 2.2 and 3.2 have not been deployed



Highest bandwidth



Lowest bandwidth



In a well designed network, and with correct deployment Zone 4 video clients work very well. If the network design is not adequate and clients are incorrectly deployed the users will be dissatisfied.

### 1.3 Qualified Switches

A small number of L3 switches are qualified for to support Gigabit Ethernet connected Avid Video clients. For ISIS 1.x the Cisco Catalyst 4948-10GE and Foundry FESX424. Added in ISIS 2.x is the Cisco Catalyst 4900M. These switches are tested with each software release.





With ISIS the release of v2.4 (AUGUST 2011) The new catalyst C4948E (already deployed with ISIS5000) receive formal approval, and also the FESX624 which is a direct IPv6 capable replacement of the discontinued (by Brocade/Foundry) FESX 424

The Avid ISIS 7000 Ethernet Switch Reference Guide available at:  
[http://avid.force.com/pkb/articles/en\\_US/Compatibility/en348609](http://avid.force.com/pkb/articles/en_US/Compatibility/en348609)

These documents describe the base configurations and network example and provides sample files configuration files as starting points to customize the configuration.

Note FESX series switches do not support the use of Priority 7 in combination with 802.1Q VLAN tags and medium/high resolution ISIS clients. This change takes effect from Code version 2.4

The workaround is to use all ports as priority 0:

```
qd 1 896 0          interface ethernet 1
qd 2 896 0          no flow-control
...                !
qd 24 896 0        interface ethernet 2
                   no flow-control
                   !
                   .....
                   interface ethernet 24
                   no flow-control
```

Alternatively the `priority ignore-8021p` command must be used on the 10G interfaces as below.

```
interface ethernet 48
  port-name TO HAVE A NAME
  no flow-control
  priority 7
```

```
interface ethernet 49
  priority ignore-8021p
  priority 7
```

### 1.3.1 Using Fast Ethernet

Gigabit Ethernet is the preferred connection for all devices, however data clients, such as Capture Manager and Control Air can connect at Fast Ethernet, as the LAN data messaging is low bandwidth and less time critical (time critical application are likely to be RS422 serial connected). Also the Avid Lo-Res Encoder has a Fast Ethernet connection, see section 1.15 for more information.

**NOTE:** Since 2008 the need to consider anything less Gigabit Ethernet in a professional broadcast environment has become almost non-existent.



### 1.3.2 Using SMC 8724ML3

Very small ISIS solutions can use the 24 Port SMC 8742 Layer 3 Gigabit Ethernet switch. This device has only been approved for data clients, not video clients. While this switch does have two 10 Gigabit Ethernet ports, using these to connect with ISIS is not supported. The typical use of this device is to connect for example a Capture Manager on the ISIS left VLAN to an AirSpeed on the right VLAN, plus give a data connection to the corporate network to receive software updates.

Note: Avid stopped selling this switch in 2009.

## 1.4 Approved Switches

Whereas qualified switches are sold by Avid, customer funded testing has been performed on many switches, some of which have been approved for direct connection to ISIS. Configuration using approved switches must be agreed with Avid Network consultants.

### 1.4.1 Cisco 6500

This switch is approved in certain configurations only. Approved 6500 Components:

WS-SUP720-3B	Catalyst 6500/Cisco 7600 Supervisor 720 Fabric MSFC3 PFC3B
WS-X6748-SFP	Catalyst 6500 48-port GigE Mod: fabric-enabled (Req. SFPs)
WS-X6748-GE-TX	Cat6500 48-port 10/100/1000 GE Mod: fabric enabled, RJ-45
WS-X6704-10GE	Cat6500 4-port 10 Gigabit Ethernet Module (req. XENPAKs)
WS-X6708-10GE	Cat6500 8-port 10 Gigabit Ethernet Module (req. X2 modules) ) restrictions apply see below
XENPAK-10GB-LR	10GBASE-LR XENPAK module
XENPAK-10GB-SR	10GBASE-SR XENPAK module



#### NOT APPROVED FAILED TESTING

**WS-X6548-GE-45AF**

**WS-X6148-GE-45AF**

**Gigabit links between switches**

**Note 1.4.1a:** WS-SUP720-3B versus WS-SUP720-3BXL The XL model has more on board memory for routes, netflow, ACLs etc, and we don't use these facilities so it should not be an issue. It may be that customer already uses WS-SUP720-3BXL and has on-site spares so it is more convenient.

**Note1.4.1b:** WS-X6724- xxx is also acceptable as it has the same architecture as the WX-X6748-xxx, but it has fewer ports and the optical version uses the larger GBIC modules.

**Note1.4.1c:** In Q1 2008 Cisco introduced some new supervisor modules with integrated 10G uplinks

Model numbers VS-S720-10G-3C and VS-S720-10G-3CXL.





The VS-SUP720-10G-3C uses a similar switch-fabric to the WS-SUP720 models, but features a new generation of PFC with -3C & 3CXL models.

The VS products are more capable and add some additional capabilities such as Virtual Switch Support (VSS) and Multi-chassis Etherchannel (MEC), Plus some on-board 10G ports.

If the solution did not use these extra functions, and only connects Avid devices to approved interface cards, it is very likely to work just fine. However as this product has not been tested by Avid, and at the time of writing, no commitment of suitability can be made by Avid.

#### **1.4.2 Use of WS-X6708-10G-3C**

Avid recommend the used of the WS-X6704-10GE interface card which has a non blocking architecture.

The used of the Cisco WS-X6708-10G interface card in a Cisco 6500 is permitted with some restrictions, as outline below:

The 6708 is oversubscribed by design. It is a blade that has 80 Gb/sec full duplex connectivity, but only has 40 Gb/sec backplane capability, with only ports 1, 2, 5 and 6 can be run at line rate to realize the full backplane bandwidth. because of the way the board has been designed to share 40 Gb/sec of bandwidth across the 8 ports. For example, if you try run ports 1,2,3 and 4 at line rate/full duplex you will see ~21% packet loss using 1518 byte frames. If you try to run all 8 ports at line rate full duplex you will incur 50.3% packet loss. There are some data flows that show greater than 40 Gb/sec within the 6708 blade. In those cases, port pairings are such that the data is not hitting the backplane; rather it is passed port to port via the Distributed Forwarding Card (DFC).

ISIS 7000 should only be connected to ports 1, 2, 5 and 6. Connecting ISIS to ports 3, 4, 7 and 8 is not supported, but these ports could be used for onward connections. Using these contended ports to connect a Catalyst 4948 should be acceptable, but it might result in reduced throughput in extreme conditions, in most situations this will not be a factor. Port 1 contend with port 3, port 2 contend with port 4, port 5 contend with port 7 and port 4 contend with Port 8.

In practice, cascading C4948 from the same card as ISIS connections can work fine because the direction of data will not cause oversubscription. However the precise nature of which ports are used and how is critical. Hence connecting ISIS to ports 1 & 3 and cascaded C4948 to port 2 & 4 is BAD because port 1 & 3 are 2:1 contended, but connecting ISIS to ports 1 and 2 and cascaded C4948 to port 3 & 4 is OK because the direction of traffic does not cause backplane contention.

When used in the corporate network (Zone 4) on links between the distribution and access layer switches the contended nature of the card should not have a negative impact on ISIS clients, providing BW demand does not exceed available capacity.

### 1.4.3 Use of WS-X6716-10G-3C

The WS-X6716-10G interface card has not tested by Avid. This card is a 4:1 oversubscribed interface card which can be used in dedicated 4 port mode. When used in the oversubscribed mode it would not be suitable for direct connection to ISIS as part of the Border Switch. When used in DEDICATED mode its capabilities exceed the WS-X6704-10GE and hence should be suitable for direct connection to ISIS as part of the Border Switch.

### 1.4.4 SUP 2T and 6800 /6900 series I/O module

These new products were release by Cisco in Q2/2011 and have not been tested by avid at (RELEASE 1.9 of DOCUMENT AUGUST 2011.

The new 6908 10G module is a 1:1 8 port device taking full advantage of the new 80Gbps slots with 256MB buffer per port. THIS PRODUCT IS LIKLEY TO BE WELL SUITED BUT HAS NOT BEEN TESTED

The 6800 series I/O modules, appear to be similar than the 6700 equivalents when it comes to buffers! In fact this is the same hardware but with a DFC4 daughter board.

The 6816 10G module is still 4:1 and uses only 40G of the 80G backplane access. THIS PRODUCT IS LIKLEY TO BE WELL SUITED BUT HAS NOT BEEN TESTED

The 6848 (/24) again uses the only 40G (20G) of the 80G access, which is not unreasonable, but they have the same 1.17MB per port buffer as the 6748.. THIS PRODUCT IS LIKLEY TO BE WELL SUITED BUT HAS NOT BEEN TESTED

There is also a the 6816 10G copper module. THIS PRODUCT IS LIKLEY TO BE WELL SUITED BUT HAS NOT BEEN TESTED

The 6816 modules should be fine for cascading down to C4948 edge access devices that are moderately loaded, and for UHRC connections, but not well suited for direct connection with ISIS unless the port group is in PERFORMANCE mode, or the bandwidth loadings are relatively low. THIS PRODUCT IS LIKLEY TO BE WELL SUITED BUT HAS NOT BEEN TESTED.



NOTE: the WS-6708-10G modules cannot be deployed with a SUP 2T. Any module with a DFC3 board cannot be used with SUP 2T

### 1.4.5 Cisco Catalyst 4500 Classic

This switch is approved in certain configurations only

WS-X4516+10GE	Cisco Catalyst 4500 Supervisor Engine V-10GE, 2 x 10 Gigabit Ethernet, console RJ-45
---------------	--



WS-X4013+10GE	Cisco Catalyst 4500 Supervisor Engine II-Plus, 2 x 10 Gigabit Ethernet, console RJ-45
WS-X4506-GB-T	Cisco Catalyst 4500 Gigabit Ethernet Module, 6 ports 10/100/1000 802.3af PoE or 1000BASE-X (SFP)
WS-X4548-GB-RJ45	RJ45—Cisco Catalyst 4500 Enhanced 48-Port 10/100/1000 Module (RJ-45) <b>FOR Interplay Production ASSIST BROWSE RESOLUTION CLIENTS ONLY</b>
X2-10GB-LR	10GBASE-LR X2 module
X2-10GB-SR	10GBASE-SR X2 module

### 1.4.6 Cisco Catalyst 4500 Enhanced

The Catalyst 4500-E has not been tested by Avid at the time of writing this document. This models feature different chassis, supervisor card and interface cards to those in the 4500 Classic. However the 4900M which is now qualified is based on the same Supervisor 6 architecture, as 4500-E but there is no direct correlation on Enhanced WX-X46xx I/O modules.

WS-X4606-X2-E	Catalyst 4500 E-Series 6-Port 10GE (X2) 1.251 Contended
WS-X4624-SFP-E	Catalyst 4500 E-Series 24-Port GE (SFP)
WS-X4648-RJ45V-E	Catalyst 4500 E-Series 48-Port PoE 10/100/1000(RJ45)

Due to the heritage of this device and the improvements in the Supervisor 6 architecture this is considered a safe, even though not officially approved.

Testing by Cisco Engineers against Avid specified test profiles and subsequent project deployment as a Zone 4 edge device, has shown that using WS-X4648-RJ45V-E in combination with the Supervisor 6 is suitable for ISIS 2.x Gigabit Ethernet clients using DNxHD resolutions editors based on a 2:1 oversubscription profile of 512KB chunks.



		2:1 OverSubscription 512KB chunk size	
		Uncontended	Contended
Tested	Single 10GE as Ingress	24	48
	Two 10GE as Ingress	24	26
Extrapolation	Single 10GE as Ingress	All	All
	Two 10GE as Ingress	All	26

A 7 slot chassis with SUP6E and the WS-X4648 Gigabit Ethernet interface card was used. Default queue limits were unchanged.

Use of additional 10Gigabit Ethernet interface cards is not supported.

The Catalyst 4500-E is not approved for direct connection to ISIS or 10G UHRC clients.



**When using 10 slot 4510E Not all slots have the fast 24Gbps backplane.**

Bandwidth Per Line Card Slot using Supervisor 6-E:

Up to 24 Gbps on slots 1-4 & 7; 6 Gbps only on slots 8-10.

**This will significantly impact performance of high speed cards used in a low speed slot.**

### 1.4.7 Cisco Nexus 7000

The Cisco Nexus 7000 is an extremely capable switch is well suited to ISIS video clients. It was tested with ISIS 7000 V2 by Avid in February 2011 and has received approval status with selected I/O modules, described below, only. It will appear as APPROVED in the qualified switch guide updated and part of the ISIS 2.4 point release due in August 2011.

1 Nexus 7000 with 48 port 1-Gb module (copper) N7K-M148GT-11,  
1 32 port 10-Gb module (optical) N7K-M132XP-12;

BIOS 3.19.0, Kickstart 4.2(4), System 4.2(4), CMP BIOS 02.01.05 CMP  
Image 4.2(1)

- The software version of the NEXUS 7000 tested was System 4.2(4), current Version at March 2011 is System 5.1(2).



The N7K-M108X2-12L uncontented 8 port 10G module is also suitable but has not been explicitly tested.

The N7K-M148GS-11 48 port optical module is also suitable but has not been explicitly tested.

The 6MB per port buffer on the 48 port Gigabit Ethernet interface card mean this module will easily capable of supporting an Avid editing clients.

The capability of the 10G interface cards and fabric, exceed those of the approved Catalyst 6500 10G interface cards (described elsewhere in this document).

A Nexus 7000 is a core grade switch and directly connecting ISIS with 10G connections (or any access level device) to the primary core contravenes best practice design. When Nexus 7000 is deployed as a local/production core/distribution direct connection from ISIS with 10G connections is viable.

Follow-on testing in July 2011 confirmed the suitability of Nexus 7000 to be used with aggregate 10G links when connecting with ISIS 7000.

The Port channel must be configured on the 10G interfaces as `channel group NN mode on`

this will show in the config file as `channel group NN`  
**mode on** not shown



Note: 10G ports on the 32 port 10-Gb module (optical) N7K-M132XP-12 which connect directly to ISIS should be used in performance mode.

### 1.4.9 Foundry/Brocade FESX 448/648

This switch is similar to the FESX 424, and can be used in place of FESX 424 because it provides additional buffering.

Note FESX/SUPERX series switches do not support the use of Priority 7 in combination with VLAN tags and medium/high resolution ISIS clients. This change takes effect from Code version 2.4

The workaround is to use all ports as priority 0:

```

qd 1 896 0          interface ethernet 1
qd 2 896 0          no flow-control
...                !
qd 24 896 0         interface ethernet 2
                    no flow-control
                    !
                    .....
                    interface ethernet 24

```



Alternatively the `priority ignore-8021p` command must be used on the 10G interfaces as below.

```
interface ethernet 48
  port-name TO HAVE A NAME
  no flow-control
  priority 7
```

```
interface ethernet 49
  priority ignore-8021p
  priority 7
```

Note: In AUG 2010 Brocade announced the End of Sale of the FESX 424 as part of a product consolidation. End of Support is MAR 2016. The existing IPv6 capable products, ESX 624/648, are a direct replacement based on the same chipset, but with some additional functions, which can be enabled by a license string. The FESX 624 was tested by Avid in FEB 2011 and found to provide equivalent support for Avid ISIS clients.

Brocade FESXv4 series

- EOL Notification Date AUG/18/2010
- Last Time Order (LTO) Final, Date FEB/28/2011
- EOL Last Ship Date (LCS) MAR/31/2011
- End of Support (EOS) Date MAR/31/2016

#### 1.4.10 Foundry/Brocade SuperX

This switch is the same family as the FESX (in the same way that the Cisco Catalyst 4948 and Catalyst 4500 are sibling products). This product has been tested in the 8 Slot version and may be used with the dual port 10G card and the 24 port SFP or 10/100/1000 interface cards in certain configurations only The 16 Slot version is not supported

SX-FI12GM-4-PREM	FastIron SuperX Management-1 module with 12 combo copper / fiber Gigabit Ethernet ports (10/100/1000 Mbps (RJ45) or Gigabit Ethernet Fiber (SFP) connectivity per port), 400 MHz processor and 256 MB SDRAM. Software includes advanced Layer 2 and full Layer 3 services (BGP-4, OSPF, RIP, VRRP, VRRPE, DVMP, PIM-SM and PIM-DM).
FI-SX1-4-DC	FastIron SuperX bundle with 8-slot chassis, fan tray and 1 DC power supply
SX-FI42XG	FastIron SuperX 2-port XFP 10-Gigabit Ethernet module
SX-FI424C	FastIron SuperX 24-port 10/100/1000 Ethernet module

The SX800 or Super X is equally suitable; the difference is that the SX has redundant management cards without additional Gigabit Ethernet ports, whereas the SuperX has a single management card with 12 Gigabit Ethernet ports.

## ALSO SUPPORTED

SX-FI424HF	FastIron SuperX 24-port 100/1000 Combo Fiber Ethernet module
------------	--

Note FESX/SUPERX series switches do not support the use of Priority 7 in combination with 802.1Q VLAN tags and medium/high resolution ISIS clients. This change takes effect from Code version 2.4

The workaround is to use all ports as priority 0:

```

qd 1 896 0          interface ethernet 1
qd 2 896 0          no flow-control
...                !
qd 24 896 0         interface ethernet 2
                    no flow-control
                    !
                    .....
                    interface ethernet 24
                    no flow-control
  
```

**NOTE: When using 10G links for cascaded switches or UHRC clients the QD setting must be correctly set.**



### 1.4.10.1 Setting for Queues on SX

With SX switch the QD setting for the 10G egress port from Core to edge must be set correctly. The default setting will restrict performance.

A QD setting of QD <m/p> 4095 and corresponding QD ,m/p> 4095 0 must be set on these ports. In fact it is recommended to set this on ALL 10G inter-switch ports that will carry ISIS traffic. And all 10G port which connect to ISIS.

It should also be applied on ports which connect UHRC clients but there appears to be little benefit to having a value above 2048. However as each 10G port is its own Packet Processor it will be simpler to apply QD 0 QD <m/p> 4095 and corresponding QD ,m/p> 4095 0 to all 10 G ports

All ports should use priority 0 not priority 7

### 1.4.11 Foundry/Brocade Big Iron RX

The BigIron RX switch was added as a solution In January 2008 and this is reflected in version 1.3 of the Avid ISIS 7000 Ethernet Switch Reference Guide available at, Current version : <http://avid.custkb.com/avid/app/selfservice/search.jsp?DocId=348609>

The RX-8 Chassis was evaluated with:  
 RX-BI-MR Management Module





RX-BI-SFM3 Fabric Module  
RX-BI-24C  
RX-BI24F  
RX-BI-4XG

Testing included 10G Link aggregation to ISIS and 1G Link aggregation with a Cisco Catalyst 6500 acting as “Zone4” switch.

This type of switch is likely to be used in large scale deployments only.

#### **1.4.12 Arista Networks 7048**

Arista 7048T switch with 48 1-Gb ports and 4 10-Gb ports, software image 4.8.6

The Arista 7048T switch has successfully passed all tests. (Tested JUL 2012)

Testing was completed on an ISIS 7000 setup with all 1- Gb clients. Only 1-Gb clients were used to maximize oversubscription. Since ISIS 7000 is more strenuous than ISIS 5000 then the Arista 7048T should work in an ISIS 5000 environment.

#### **1.4.13 Cisco Catalyst C4500-X**

Cisco 4500X switch with 32 10/1-Gb ports with IOS 03.03.00.SG & ROM 15.0(1r)SG6

The Cisco 4500X switch has successfully passed all tests. (TESTED AUG 2012)

Testing was completed on an ISIS 7000 setup with 1- Gb and 10-Gb clients. Initial testing was done with only 1-Gb clients in order to maximize oversubscription. Additional testing with 10-Gb and 1-Gb clients was also performed.

Since ISIS 7000 is more strenuous than ISIS 5000 and ISIS 2000 then the Cisco 4500X should work in both ISIS 5000 and ISIS 2000 environments.

At the time of writing (AUG 2012) the C4500-X is an approved switch for use with ISIS 7000 and ISIS 5000 and ISIS 2000. And will be added to the Approved switch guide at the next release.

At the time of writing (AUG 2012), the C4500-X is an approved switch and not a qualified switch, i.e. one that is supplied by Avid and will be tested with every major software point release as is the practice for Catalyst C4900M and C4948E. The similarities across the Catalyst C4XXX family and the superior architecture of the C4500-X vs. earlier family products negate the need for this step at this time This is subject to change should Avid resell the C4500-X.

VSS operations have not been tested and are not supported.

## **1.5 Non Approved Switches, Modules and VoIP**

Just because a switch has not been tested by Avid does not mean it will not be suitable for an ISIS client. However, if your system encounters problems Avid will be unable to support the system. Another key point to consider when the switch is deployed in Zone 4, the responsibility to administer this switch and provide necessary QoS resides with the customer not Avid.

A full analysis of the workflow and end to end dataflow may find that the chosen non-approved switch is in fact suitable. Avid Professional Services group can be engaged to provide consultancy to determine such requirements and solutions.

In the sections below some alternative products are discussed along with their suitability, or lack of suitability, depending on workflow.

### **1.5.1 Cisco Catalyst 3750**

The 3750E (also 3560E) has 10G uplink interfaces but the small 2MB buffer per 24 ports mean that this product cannot be deployed for editing (ISIS medium resolution) clients, any client using the default setting WILL FAIL! When ISIS clients are set as Low resolution, this wire-speed switch is suitable of supporting Interplay Assist/INews Instinct even at some resolutions considered HD. Also this configuration is suitable for Interplay Access (with or without Interplay Streaming server) and INews data only clients.

When the 3750E is uplinked at 1G to a 4900/4948 to multiple 1G connecting clients, this relies on the excellent dynamic buffer of the 4900/4948 to support the bursty traffic which will oversubscribe the downlink port. The quantities of concurrent ISIS clients which can be supported depend on ISIS version and video resolution.

Note: this should also apply to the Catalyst 3750X (and 3560X) which is essentially a re-packaged and cost reduced version of the Catalyst 3750E, without permanently integrated 10G electronics which must now supplied separately. But as at JAN2011 no proven experience of this product is available.

The legacy Cisco Catalyst 3750G (also 3560G) cannot successfully be directly connected to ISIS via 10G. However there are some deployed installations that cascade a 3750G from a Cisco Catalyst 4948/4500 or 4900M via a 1G uplink to multiple 1G connecting clients. This relies on the excellent dynamic buffer of the 4900/4948 to support the bursty traffic which will oversubscribe the downlink port. The quantity of concurrent ISIS clients which can be supported depend on ISIS version and video resolution.

The 3750G is suitable for Interplay Access (with or without Interplay Streaming server) and INews data only clients.

### **1.5.2 Cisco Nexus 5500/5000/2000**

The Nexus 5500/5000 is a layer 2 switch with small buffers, hence unsuitable for deployment of Gigabit Ethernet connected Avid editing clients, it has not been tested by Avid and its status is not approved. The same applies to the Nexus 2000 module. When ISIS clients are set



as Low resolution, this wire-speed switch should be able to supporting Interplay Assist but would be a very expensive solution for this task.

Note: The one customer that unsuccessfully attempted to deploy this switch against the advice of Avid, replaced it with the qualified Cisco Catalyst 4900M.

The Nexus 5000 might be suitable to connecting 10G UHRC clients as Zone 2 devices, but it has not been tested by Avid in this way and its status is not approved.

This switch is unsuitable for use with ISIS because NEXUS 5000/5500 series products use ingress buffering which is incompatible for AVID ISIS traffic toward 1G clients which inherently oversubscribes the egress path. ISIS 7000 versions up to 2.4 with default install MEDIUM resolution have a 2.0x oversubscription profile. ISIS 5000 has a 1.5x oversubscription profiles, and this was applied to ISIS 7000 clients from V3.5. Whereas other switches use shared packet buffer (C49xx) or dedicate edge port buffer (Nexus 7000), The Nexus 5000 has approx. 480KB ingress buffer per port and the Nexus 5500 has approx. 660KB buffer per port with defined pools for Ingress and egress. This quantity of ingress buffer is not sufficient for supporting multiple concurrent 1G clients sets as the default MEDIUM resolution. Changing the client setting to LOW resolution would mitigate these limitations but this is only really suitable for browse style client, not editors or servers. While an oversubscription profile of 1.5x will reduce the chance of packet loss in Nexus 5000 ingress buffers, the probability remains too high in medium to high activity systems.

#### UPDATE MARCH 2012

Joint testing by Cisco and Avid, using a combination of Nexus 5000 and 2248 TPE showed a marked improvement for this platform family, however the shallow ingress buffer architecture remains a concern, as it is a weak point when dealing with bursty traffic from the ISIS, and it is easy to cause packet drops under moderate loads toward 1G clients. A single Gigabit ISIS client (default Medium resolution) connected directly to a 1G port on the Nexus 5548 will encounter packet drops when executing PATHDIAG from ISIS. While the Nexus 5500 is appropriate to traditional data-centre LAN environments, the Nexus 5000 architecture continues to be inherently unsuitable for use with ISIS. Customers should consider the Nexus 7000 platform instead which now has smaller form-factor 4 slot chassis.

#### **1.5.4 4908 10G module for Cisco Catalyst 4900M**

This is a 2:1 contended I/O module and has not been tested by Avid. It should not be connected directly to Avid ISIS. However, it should be suitable for connecting low demand cascaded switches (e.g.494810-GE) or 10G UHRC clients where the average maximum BW demand does not exceed 5Gbps. Use of this I/O module is not explicitly supported, nevertheless implemented correctly it should be well suited, and has been successfully deployed in several customers supporting subordinate C4948 switches.

#### **1.5.5 Juniper EX3200 and EX4200**

These two switches from Juniper are similar in characteristics to Cisco 3750E with respect to ISIS video traffic. When uplinked by 10G interfaces but the small buffer per 24 ports mean that this product cannot be deployed for editing (ISIS medium resolution) clients. When ISIS clients are set as ISIS Low resolution, this wire-speed switch should be well suited to supporting Zone 4 Interplay Assist even at some resolution considered HD. These Juniper



switches have not been tested by Avid in this way and its status is not approved. The responsibility to provide required QoS in Zone 4 reside with the customer.

Note: Some basic testing (with a single clients only) in MAY 2010 supported the above assessment of suitability and limitations.

### 1.5.6 Brocade/Foundry NetIron MLX

This switch has many similarities to the Foundry BigIron RX which has been approved for ISIS. It is a chassis based switch but uses some similar and some different I/O modules to The Big Iron. In Q4 2009, Avid observed some Lab testing of an ISIS system directly connected to a MLX 4 switch with 10G and 1G clients. This Brocade switch has not been tested by Avid in this way and its status is not approved.

The NETIRON MLX with V 4.x software, appeared capable of supporting Gigabit Ethernet and 10Gigabit Ethernet clients and bears many similarities with the BIGIRON RX, in fact in some areas MLX appears to exceed the capabilities of the RX, the these aspects are not well defined in publically available documentation.

The Link aggregation load balancing algorithms available within the MLX are unsuitable for use with ISIS, and further development from Brocade would be required to resolve this limitation.

#### 1.5.6.1 Brocade/Foundry NetIron MLXe

As part of a customer funded consultancy project (MAR 2011), the MLXe platform with V5.x s/w was extensively tested (alongside and SX1600 setup) with Zone 2.1 and Zone 3.1 10G clients, and link aggregation. This has not been tested by Engineering in Burlington and is awaiting implicit approval based on the external test report).

This switch is considered suitable by the author but is not explicitly approved.

NI-MLX-1Gx20-SFP	NetIron MLX Series 20-port FE/GE (100/1000) module with IPv4/IPv6/MPLS hardware support - requires SFP optics. Note: Copper SFPs are supported at 1000Mbps only
E1MG-SX-OM	1000Base-SX SFP optic, MMF, LC connector, Optical Monitoring Capable
NI-MLX-1Gx20-GC	NetIron MLX Series 20-port 10/100/1000 copper module with IPv4/IPv6/MPLS hardware support
NI-MLX-10GX8-D	MLX 8-port 10-GbE (D) SFPP module



**NOTE The earlier MLX platform with Version 4.x s/w is not Supported as there are known issues with 10 G link aggregation to ISIS**

#### 1.5.6.2 Setting Queues on MLXe

Change queue 0 to 8192KB using the command:

```
qos queue-type 0 max-queue-size 8192
```

to ensure sufficient egress buffers for “remote” 10G and 1G clients.

Each traffic manager supplies at the ingress side 8 priority queues that are mapped to each egress port (max. 8000 queues/TM). The default value for each queue depth is 1MB.

If there is congestion at the egress port, the traffic manager responsible for the egress port (on the egress module) will make sure that packets are being sent according to their priority.

Module Per /TM of Memory	# of Traffic Total Amount	Managers (TM) per module		Amount of memory
8x10Gb/s	2	512MB	1GB	
24x1Gb/s SFP	1	512MB	512MB	
20x1Gb/s Cu	1	128MB	128MB	

Since packets are actually buffered on the ingress side this method is called Virtual Output Queuing.

Queues can only be adjusted globally with the following command (e.g. for priority queue 7):

```
qos queue-type 7 max-queue-size 8192
```

This will increase the size of queue 7 from default 1024KB to 8192KB, the maximum value is 32768KB, which not necessary for ISIS traffic.

### 1.5.7 Switch Buffering architectures and limitations

The qualified switches, (Cisco Catalyst 49xx and Brocade/Foundry FESX) were chosen because of their excellent dynamic buffering characteristics for 10G to 1G operations. These benefits are also apparent in the sibling products (Cisco Catalyst 4500 and Foundry FESX 448/Super X). These qualities are particularly important for a system that will routinely have multiple Gigabit Ethernet connected servers ( ISIS SERVER BLADES) sending data to a single Gigabit Ethernet port on an external switch.

The 6500/6748 architecture has more memory than a 4500/4948 but it is statically assigned to individual ports, and unused buffer on one port cannot be allocated to a port which needs more than its normal assignment. This limitation is exacerbated when the available buffer is further divided by the use of QoS queuing strategies. This restricts the design options for this in terms of the number of streams which can be supported on an individual 6500 series Gigabit Ethernet port. While individual SD clients can be supported, cascading an access



switch for SD video streams is not viable, but a limited number of low resolution streams can be supported on a cascaded switch. Designs using this architecture must be approved by Avid.

The 3750G series of switch cannot be directly connected to ISIS via 10G due to buffering limitations, but it can be deployed as cascaded access switch supporting SD or low resolution clients. Designs using this architecture must be approved by Avid.

### 1.5.8 Inline VoIP device

Inline VoIP devices are not supported for Avid ISIS clients that play real time video. The primary reasons for this are:

- (i) most IP Phones to date have been 100Base-T Fast Ethernet devices and this is not a suitable connection for real time video editing. Only recently are 1000Base-T Gigabit Ethernet devices beginning to reach the market.
- (ii) The performance 100Base-T Fast Ethernet devices is a substandard user experience compared to 1000Base-T Gigabit Ethernet devices, even with low video resolutions.
- (iii) that the Cisco 6500 series WS-X6148-GE and WS-X6548-GE interface cards (PoE capable) that have been tested, do not have sufficient buffering to support the 2:1 oversubscription required by an ISIS Medium Resolution editor client.
- (iv) Even with Gigabit Ethernet connecting devices, the QoS profiles deployed on the port may have a negative impact on the editing client plus the configuration complexity will increase on the switch.



Note: the WS-X6748-GE Interface card does not support PoE.



Note: While the Cisco Catalyst 4500 series switches have sufficient buffering to deal with VoIP and have interface that can support PoE, the different families of interface cards provide varying levels of buffering depending on Supervisor model.



Note: While the Foundry FESX and Super X series switches have sufficient buffering to deal with VoIP and have interfaces that can support PoE, the configuration requirements of the edge device are critical as the default queue depths are not suitable.

Note this restriction does not apply to Avid Interplay Access clients receiving Quick time wrapped WAN browse class video from Interplay Stream Server. See section 2.8 for more information on Interplay Stream Server

### 1.6 Network Interface Card Requirements

The general recommendation is to use the Intel Pro/1000 PT/PF (PCI-E) network interface card in HP workstations such as the XW8600 or Z800. Legacy HP workstations such as such XW8200 and XW8400 required Intel Pro/1000M (PCI-X) network interface card in workstations. This high performance NIC provides the ability to adjust the descriptors or buffers to improve performance. Details of how to set up the interface card are provided in



Avid ISIS 7000 Client User's Guide v2.3

[http://avid.force.com/pkb/articles/en\\_US/User\\_Guide/en389015](http://avid.force.com/pkb/articles/en_US/User_Guide/en389015)

Avid ISIS Client User's Guide v 4.0

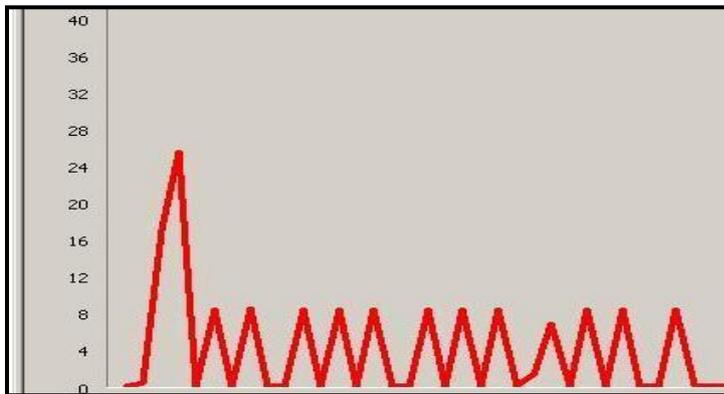
[http://avid.force.com/pkb/articles/en\\_US/user\\_guide/Avid-ISIS-Client-User-s-Guide-v4-0?popup=true&NewLang=en&DocType=1082](http://avid.force.com/pkb/articles/en_US/user_guide/Avid-ISIS-Client-User-s-Guide-v4-0?popup=true&NewLang=en&DocType=1082)

### 1.6.1 Using Fast Ethernet

Gigabit Ethernet is the recommended connection for all ISIS 7000 clients. It is possible to connect and Avid Interplay Assist client using MPEG II browse resolution via Fast Ethernet. While steady state DV25 is 4MB/s and Fast Ethernet provides 12MB/S, media may not play reliably on a Fast Ethernet connection, as it will cause delays when filling buffers which typically require 4-5 x the steady state rate during play start, i.e. 16-20MB/S. During continuous play the data profile will peak up to 8MB/s.



**Use of any Fast Ethernet based solution needs the agreement and approval of Avid network consultants.** Every point in the path between ISIS and Client must be fully understood to ensure sufficient buffering exists



DV25 1V 0A (4 MB/s)  
Medium Resolution Selected  
Read ahead Enabled (OS=1)



DV25 1V 0A (4 MB/s)  
 Low Resolution Selected  
 Read ahead Disabled (OS=0)

### 1.6.2 When not to use the Intel PRO/1000M or Intel Pro/1000P NIC

Whilst the generally recommendation is to use this card there are some exceptions. The HP XW4300 and XW4400 should not use this card, instead using the on-board NIC. Use of Intel NIC in XW8600 is now optional for 256K chunk size in ISIS 1.x but required for 512KB chunk size in ISIS 2.0.

See knowledge base articles:

- What Network Interface Controller (NIC's) are Supported in Client Systems on Avid ISIS 7000?

[http://avid.force.com/pkb/articles/en\\_US/Troubleshooting/en244563](http://avid.force.com/pkb/articles/en_US/Troubleshooting/en244563)

- Performance enhancement for Broadcom Nic on ISIS

[http://avid.force.com/pkb/articles/en\\_US/Compatibility/en266865](http://avid.force.com/pkb/articles/en_US/Compatibility/en266865)

In Q2 2008, the XW8600 has been qualified for use with Media Composer & NewsCutter. An additional Intel Pro series card is now optional rather than highly recommended as for earlier models. The onboard NIC is a Broadcom 5755 dual controller and is now sufficient for most applications using ISIS 1.x. The registry modification to enhance performance for Broadcom NIC on ISIS (described above) should be applied.

There is one exception. If you have an HP XW8600 Vista 64 Media Composer Adrenaline client then you need to use the Intel Pro 1000 PT to connect to either ISIS or MediaNet Ethernet attached. Using the onboard Broadcom with the Adrenaline in this scenario causes data under runs while doing a digital cut or output to a playback monitor.

The following important limitations are now in effect for ISIS 2.x. This information will be added to the 2.0.1 readme shortly. Note that if you are utilizing i2000 ISBs, you must select 512k chunk size and are therefore subject to this limitation.



The following important limitations are now in effect for ISIS 2.x. This information will be added to the 2.0.1 readme shortly. Note that if you are utilizing i2000 ISBs, you must select 512k chunk size and are therefore subject to this limitation.

### 1.6.3 HP xw8600 & Z800 & Z400 Workstations and Broadcom Ethernet Connections

Avid ISIS 7000 Clients using a built-in network port utilizing the Broadcom® chipset and a Storage Group with a 512k chunk size, are limited to resolutions that draw 16 MB/s (50 Mb/s) or less. These clients include; laptops, Avid Interplay Assist, Avid iNews Instinct, and the HP xw8600 workstation using the built-in network port.

The built-in network port can be used for Ethernet connections to the Avid ISIS 7000 infrastructure when:

- using workspaces with 256 KB chunk sizes
- using workspaces with 512 KB chunk sizes and bandwidths of 16 MB/s or less

*NOTE For example, you can use the built-in Broadcom chipset to run two streams of DV 50 or DNxHD 36. Bandwidths are listed by resolution and number of streams on the Avid Knowledge Base. Search the Avid Knowledge Base for the Avid ISIS 7000 Performance Guide.*

You are required to use the Intel Pro 1000 PT for your network connection in the xw8600 when:

- using workspaces with 512 KB chunk sizes and your xw8600 workstation using bandwidths higher than 16 MB/s
- using a Windows Vista, 64-bit client connected to an Avid Adrenaline or Mojo.



**See the most recent readme for precise information**

### 1.6.4 Setting descriptors

As discussed elsewhere in this document, descriptors are buffers, needed for in the case of ISIS traffic for re-assembling fragmented UDP datagrams.

The example below shows the differences when an Integrated Intel NIC (Intel 82566DM Gigabit Network Connection) was tested to see what differences were experienced when the NIC default descriptor allocations of 512TX and 256RX versus 1024TX/1024 RX against Path DIAG unlimited.





avid.com



Figure 1 default descriptor allocations of 512TX and 256RX

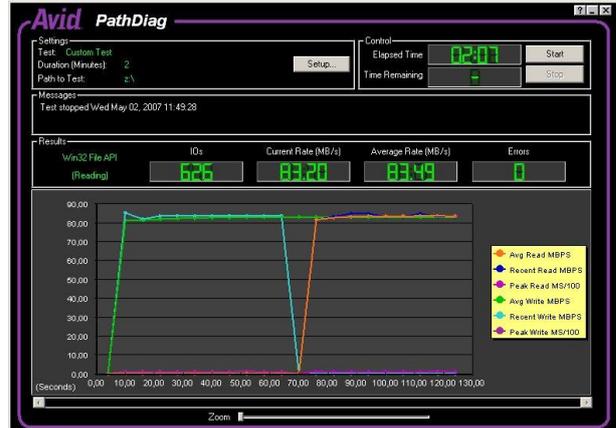


Figure 2 descriptor allocations of 1024TX and 1024RX

The Device under test showed approx 80MBWrite and 80MB Read when running and Unlimited PATHDIAG of with 8192 transfer size. A lower read performance of approximately 50MB/s would be apparent if the device was set for Low Resolution in the ISIS 7000 Client. Increasing read descriptors to 1024 had a minor positive effect on the throughput, but this is not beneficial to an Interplay Assist client at MPEG II resolution.

**Guidelines:**

- When the ISIS client is set for low resolution there is minimal benefit in increasing the descriptors.
- When the When the ISIS client is set for medium or high resolution there is significant benefit in increasing the descriptors to the maximum setting.

On the Intel platform descriptors are set using the Network Connection Properties.

For a Broadcom NIC the settings this is achieve by adjusting the registry.

- Performance enhancement for Broadcom NIC on ISIS

[http://avid.force.com/pkb/articles/en\\_US/Compatibility/en266865](http://avid.force.com/pkb/articles/en_US/Compatibility/en266865)

**1.6.5 10G network interfaces for Ultra High Resolution Clients**

When Avid introduced Ultra High Resolution Clients in Q4/2009 the preferred deployment was Zone 1 only in an Avid Unity compatible single channel 10 GbE PCI-e network interface card (CHELSIO) with short range (SR) integrated optics. For use with xw8600 workstations and dual quad-core xw8400s.

In Q3/2010 introduced the Avid ISIS 7000 compatible single channel 10 GbE PCI-e network interface card (MYRICOM) with short range (SR) integrated optics. For use with HP Z400 or Z800 workstations, with Windows XP32, Windows Vista64, or Windows7. Also for use with dual quad-core MacPro workstations running Snow Leopard.



### 1.6.7 LAN on Motherboard

MOBO or LoM implementations are often suitable. Experience has proved that those based on the Intel chipset are suitable for use in Zone 4 clients running Interplay Assist or iNews Instinct. In fact the Integrated Intel NIC on the HP Z200 workstation is approved for Media Composer/NewsCutter/Assist/Instinct.

### 1.6.8 Intel Pro 1000 CT gigabit adapter

Testing in July 2011 (EMEA) has indicated that the Intel Pro 1000 CT desktop Gigabit adapter is suitable for use in Zone 4 clients running Interplay Assist or iNews Instinct. Just as with LoM this NIC is not explicitly tested or approved by Avid Burlington.

### 1.6.9 Avid Slot Configuration guide

The Avid Configuration Guidelines and Slot Configurations documents available at URL: [http://avid.force.com/pkb/articles/en\\_US/User\\_Guide/en269631](http://avid.force.com/pkb/articles/en_US/User_Guide/en269631) provides excellent information on which Intel chipsets or interfaces cards are suitable for use with ISIS .

## 1.7 DNS

Because ISIS has dual VLANs, the use of DNS is considered essential in all but the smallest ISIS deployments, i.e. a Single ISIS Engine and single System Director. The addition of a secondary System Director makes DNS essential. When Interplay Production is added to the solution then it becomes mandatory. The DNS should contain zones for both ISIS VLANs and any other Avid-specific VLANs, such as transfer VLANs, Cluster VLANs, etc.

As useful document, “DNS for ISIS and Workgroups”, can be found at: [http://avid.force.com/pkb/articles/en\\_US/Troubleshooting/en241765](http://avid.force.com/pkb/articles/en_US/Troubleshooting/en241765)

Note: the referenced document is somewhat dated, mentioning only Windows 2000, however, the overall points are still applicable.

If a local DNS is not available, and DNS is provided via the corporate network it means that the Interplay Production/ISIS environment cannot be disconnected from the corporate network in the event of a security issue. If the link to the corporate network is broken, either due to deliberate action or a failure, the ISIS and Interplay Production will not function correctly

For optimum performance, Interplay Production/ISIS requires a DNS lookup with negligible latency, this is another reason why a local DNS is preferable.

If integration with an existing corporate DNS structure is required one viable option would be to setup DNS servers within the ISIS/Interplay Production core to serve these zones, and setup forwarders to forward all other unknown lookups to the corporate DNS servers.

### 1.7.1 DNS naming conventions

A common question regarding DNS names surrounds the use of the UNDERSCORE character.

Underscore were disallowed in hostnames since the publication of RFC952 (1985). Of course Microsoft had a different view than the UNIX world. However while a Windows DNS machine will typically allow the use of UNDERSCORE, a standards-compliant one should not.

For help choosing good hostnames, refer to RFC 1178: <http://tools.ietf.org/html/rfc1178>. While this article is a bit dated considering the growth of the Internet and IP in general, it makes some good points

Below is an extract from Wikipedia

Restrictions on valid host names  
<http://en.wikipedia.org/wiki/Hostname>

Hostnames are composed of series of [labels](#) concatenated with dots, as are all [domain names](#).<sup>[1]</sup> For example, "en.wikipedia.org" is a hostname. Each label must be between 1 and 63 characters long<sup>[2]</sup>, and the entire hostname (including the delimiting dots) has a maximum of 255 characters.

The Internet standards ([Request for Comments](#)) for protocols mandate that component hostname labels may contain only the [ASCII](#) letters 'a' through 'z' (in a case-insensitive manner), the digits '0' through '9', and the [hyphen](#) ('-'). The original specification of hostnames in [RFC 952](#), mandated that labels could not start with a digit or with a hyphen, and must not end with a hyphen. However, a subsequent specification ([RFC 1123](#)) permitted hostname labels to start with digits. No other symbols, punctuation characters, or white space are permitted.

While a hostname may not contain other characters, such as the underscore character (\_), other DNS names may contain the underscore. Systems such as [DomainKeys](#) and [service records](#) use the underscore as a means to assure that their special character is not confused with hostnames. For example, `_http._sctp.www.example.com` specifies a service pointer for an SCTP capable webserver host (www) in the domain example.com.

A notable example of non-compliance with this specification, [Microsoft Windows](#) systems often use underscores in hostnames. Since some systems will reject invalid hostnames while others will not, the use of invalid hostname characters may cause subtle problems in systems that connect to standards-based services. For example, RFC-compliant mail servers will refuse to deliver mail for MS Windows computers with names containing underscores.

### 1.8 Cable Requirements

Avid supports the following cable types for connecting an Avid ISIS system. As described in the Avid Products and Network Site Preparation Guide

Avid Products and Network Site Preparation Guide • 0130-30628-01 Rev. A • May 2011 • Created 5/16/11

Available at:

Avid Products and Network Site Preparation Guide[[373751] MAY 2011  
<http://avid.custkb.com/avid/app/selfservice/search.jsp?DocId=373751>

Cable Name	Function	Maximum Distance
Avid engine Interconnect CX-4 cable. Only available from Avid.	Connect engines. See <a href="#">Avid Engine CX-4 Interconnect Cable</a> for proper removal.	3 supported lengths at this time: 1, 3 and 5 meters
RJ45 Cat 5E or Cat 6 or Cat 6a Ethernet cable	Ethernet Avid Unity clients System Directors and clients to 1 Gb ports on an ISS Avid Interplay servers to shared storage networks Avid AirSpeed capture and playback servers to shared storage networks Avid ISIS 7000 management port to laptop	RJ45 connector 100 Meters; If using CAT5e the cable must be rated for 350 MHz for maximum length. The minimum GigE cable length for Avid network products is 6 feet or 2 meter.
Cat5 cable is <i>not</i> supported for these connections.		
<b>Avid ISIS 7000 and Avid Interplay</b> optical cable	Connects: 10-Gb port of switch to optical 10-Gb port on the Avid Unity ISIS engine. ISS 10-Gb optical port to switch port ISS 10-Gb optical port to 10-Gb Ethernet Client 10-Gb Client to 10 Gb Switch port 10-Gb Ethernet switch to 10-Gb Ethernet Switch ISS to 10-Gb adapter in Move/Copy service	The maximum length for 10 Gb Ethernet cable is defined by the core diameter (measured in microns) and modal bandwidth (in units of MHz*km). Avid supports multi-mode fiber cable using 850 nm transceivers (short distances). Specifications for these cables can be found in the ISO 11801 structured cabling document. <b>MMF 62.5 micron cable Modal Bandwidth of:</b> (Overfilled Launch (OFL) Bandwidth, typical of OM1 cable) • 160 MHz*km at 26 meters • 200 MHz*km at 33 meters <b>MMF 50 micron cable Modal Bandwidth of:</b> • 500 MHz*km at 82 meters (Overfilled Launch Bandwidth, typical of OM2 cable) • 2000 MHz*km at 300 meters (Effective Modal Bandwidth, typical of OM3 cable) <b>Avid supports single-mode fiber cable using 1310 nm transceivers</b> (long distances): • SMF ITU G.652.A/B 9 micron cable up to 10 km
<b>Avid ISIS 7000</b> X2 optical transceivers	Transceiver used in: Cisco 4948 and 4900M	SC connector X2 = Cisco X2-10GB-SR for MMF

		X2 = Cisco X2-10GB-LR for SMF
<b>Avid ISIS 7000</b> XFP optical transceivers	Transceiver used in: Foundry FESX424 and ISIS ISS1000	LC connector XFP = 10G-XFP-SR for MMF XFP = 10G-XFP-LR for SMF XFP = Foundry 10G-XFP-SR or Picolight XXL-SC-S45-21 for MMF XFP = Foundry 10G-XFP-LR or Bookham 10G-BASE-LR for SMF
<b>Avid ISIS 7000</b> SFP+ optical transceivers	Transceiver used in: ISIS ISS2000	LC connector • SFP+ short range (SR) Picolight PLRXPL-SC-S43-21-N or Finisar FTLX8571D3BCL or Avago AFBR-700SDZ for MMF • SFP+ long range (LR) Finisar FTLX1471D3BCL for SMF Avago AFCT-701SDZL for SMF JDSU JSH-01LWAA1 for SMF

### 1.8.1 Copper cabling for Gigabit Ethernet

In the ISIS setup guide Avid recommends that Cat6 is used wherever possible ANSI/TIA/EIA-568-B.2-1 published on June 20, 2002, [www.tiaonline.org](http://www.tiaonline.org)

The TIA Category 6A standard (TIA 568 B.2-10) was ratified by TIA TR 42 engineering committee on February 8, 2008. The TIA standard defines the requirements of four-pair balanced copper cabling to support 10G transmission for distances up to 100 meters. It represents the most advanced set of network cabling requirements specified up to 500 MHz. Category 6A is fully backward compatible with all the previous categories, including Category 6, Category 5e and Category 5.

An excellent article explaining the nebulous issues of cabling standards past present and future “De-Mystifying Cabling Specifications From 5e to 7A” by Valerie Rybinski (Updated December 2007) can be found at [http://www.siemon.com/us/white\\_papers/07-03-01-demystifying.asp](http://www.siemon.com/us/white_papers/07-03-01-demystifying.asp)

Many sites with ISIS 7000 have successfully deployed SYSTIMAX cable solutions, however this is not a specific Avid recommendation, but it is a popular choice.

Some useful articles and whitepaper references are given below:

**Category 6 Cabling: A Standards and Systems Overview.**

<http://www.generalcable.com/NR/rdonlyres/2C3E425A-E4F2-40A3-9ADA-CDBF5225CD60/0/Cat6Cabling2.pdf>





**Gigabit Ethernet in structured building cabling**

[http://www.harkis.harting.com/WebHelp/GBEthernet/WebHelp/GBEthernetGigabit\\_Ethernet\\_in\\_structured\\_building\\_cabling.htm](http://www.harkis.harting.com/WebHelp/GBEthernet/WebHelp/GBEthernetGigabit_Ethernet_in_structured_building_cabling.htm)

A very informative Category 6A whitepaper “**The End for Category 5e, Category 6 Under Fire, Long Live Category 6A**” is available from the URL below

[http://goliath.ecnext.com/coms2/gi\\_0198-349871/10GBase-T-standard-may-shift.html](http://goliath.ecnext.com/coms2/gi_0198-349871/10GBase-T-standard-may-shift.html)

or

[http://delivery.qmags.com/d/?pub=CIM&upid=11047&fl=others%2fCIM%2fCIM\\_20060801\\_Aug\\_2006.pdf](http://delivery.qmags.com/d/?pub=CIM&upid=11047&fl=others%2fCIM%2fCIM_20060801_Aug_2006.pdf) and click download depending on browser version

EN50173PL Class E Cat 6.		
<a href="http://www.broadbandutopia.com/caandcaco.html">http://www.broadbandutopia.com/caandcaco.html</a>		
Cat5e and Cat6 Comparison Category 6 Cabling System and Application Seems to undermine Cat 5e at 350MHz		

Cat 5e ANSI/TIA/EIA-568-B published on May 2001, [www.tiaonline.org](http://www.tiaonline.org) should only be used in legacy installations.

### 1.8.2 Fibre Optic cabling for 10 Gigabit Ethernet

For short haul use, Avid recommends the use of a 50uM multimode fibre with a minimum Effective Modal Bandwidth of 2000 also known as OM3 grade fibre. This allows 10 Gigabit Ethernet connections up to 300m and Gigabit Ethernet to 1000m.

Single mode fibre would allow the most flexibility for future use; alternatively 50uM Multimode Fibre with EMB of 4700 (classified as OM4 August2009) would allow 10 Gigabit Ethernet up to 550m (Corning InfiniCor eSX+ fiber).

- All multi mode fibre must conform to: TIA/EIA 492AAAC, IEC 60793-2-10 ,Type A1a.2 fiber, ISO/IEC 11801, Type OM3 fiber
- The fiber shall guarantee the following Effective Modal Bandwidth (EMB) at 850 nm, EMB ≥ 2000 MHz.km TIA/EIA 455-220, IEC 60793-1-49. [Recommended media, Corning InfiniCor SX+ fiber]
- All single mode fibre must conform to ITU-T G.652 (Categories A,B,C & D), IEC 60793-2-50 (type B1.1 & B1.3) and TIA/EIA-492CAAB. [Recommended media, Corning SMF-28e® fibre]



An excellent reference document on fibre optics is available from the Corning web site at 50 µm Fiber Q & A. Answers questions and concerns surrounding 50 µm fiber.  
<http://www.corning.com/docs/opticalfiber/wp4044.pdf>

Also See:  
 UNDERSTANDING OM1, OM2, OM3, OS1, OS2 and more!  
<http://www.fia-online.co.uk/pdf/Whites/wp0208.pdf>

Also See:  
[http://en.wikipedia.org/wiki/Multi-mode\\_optical\\_fiber](http://en.wikipedia.org/wiki/Multi-mode_optical_fiber)

Also See  
 BELDEN OM4 Optical Fiber Cabling  
[http://www.belden.com/pdfs/Techpprs/OM4\\_WP.pdf](http://www.belden.com/pdfs/Techpprs/OM4_WP.pdf)

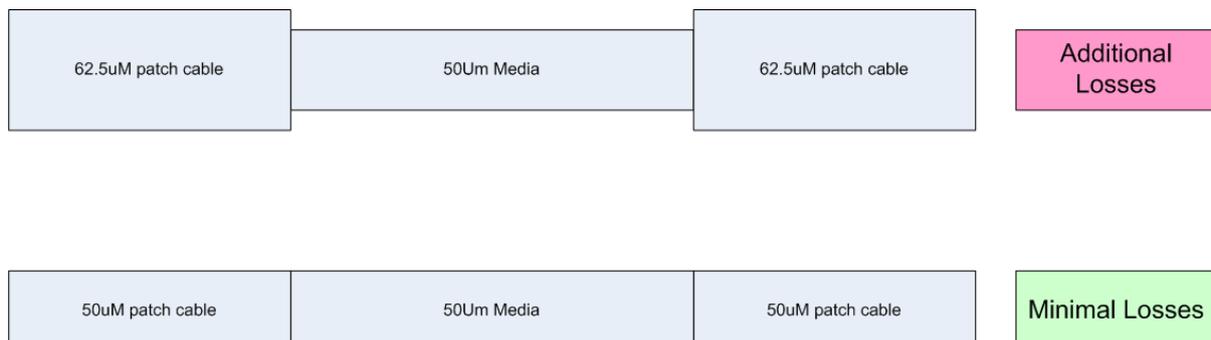
Test Requirement	OM1 – 62.5 µm	OM2 – 50 µm	OM3 – 50 µm	OM4 – 50 µm	Single Mode
OFL Bandwidth @ 850/1300 nm	200/500 MHz·km	500/500 MHz·km	1500/500 MHz·km	3500/500 MHz·km	No Requirement
EMB @ 850 nm	No requirement	No requirement	2000 MHz·km	4700 MHz·km	No Requirement
Minimum reach @ 1 Gb/s	275/550m* (850/1300nm)	550/550m* (850/1300nm)	800/550m* (850/1300nm)	1100/550m* (850/1300nm)	5000m* (1300nm)
Minimum reach @ 10 Gb/s	33m*	82m*	300m* (850nm)	550m* (850nm)	10000m* (1300nm)
Minimum reach @ 40 & 100 Gb/s	No requirement	No requirement	100m* (850nm)	125m* (850nm)	10km/40km* (1310/1550nm)

\*Note: The distances listed are industry standard minimums. Transmission distance is largely a factor of glass manufacturers' specifications and transmitter/receiver selection. Actual transmission distance may considerably exceed the distances stated. Contact Belden for current specifications.\*

OM4 is a laser-optimized, high bandwidth 50µm multimode fiber. In August of 2009, TIA/EIA approved and released 492AAAD, which defines the performance criteria for this grade of optical fiber.

### 1.8.2.1 Patch cables

Make sure you patch cable use the same size of MMF cable. When using 50uM optical media the patch cords should be the same otherwise there is a physical mismatch that introduces additional loss.



### 1.8.2.2 Labeling of Media

It can be very difficult to work out what type of fibre cable plant is installed, in an existing site, even if you can find the drawings. When implementing structured fibre cabling it is always a good idea to engrave on the panel, the media type, i.e. the SIZE and the EMB rating For example “50Um MMF EMB=2000”.

#### **Corning® InfiniCor® multimode fibers**

With superior technology and profile control, Corning® InfiniCor® multimode fibers revolutionize performance in local area networks, storage area networks and central office interconnects. InfiniCor® eSX+ and SX+ fibers, the newest offerings in the InfiniCor fiber product line, are optimized for high performance with laser-based protocols such as 10 Gigabit Ethernet.

[http://www.corning.com/opticalfiber/products/infinicor\\_fibers.aspx](http://www.corning.com/opticalfiber/products/infinicor_fibers.aspx)

**These MMF products are ideal for use with ISIS 10G links. 50uM PREFERRED**

#### **SMF-28e® fiber**

SMF-28e® fiber is the industry leader in comprehensive fiber performance. Capable of full-spectrum transmission, SMF-28e fiber is a more versatile standard single-mode fiber. The low water peak attributes combined with very low hydrogen-induced loss provide lifetime performance. SMF-28e fiber is ITU-T G.652.D-compliant and optimized for un-amplified regional, metropolitan, local access telephony and cable television networks and emerging applications.

<http://www.corning.com/opticalfiber/products/SMF-28e+fiber.aspx>

SMF 28e fiber conforms to the major optical fiber industry standards including ITU-T G.652 (Categories A,B,C & D), IEC 60793-2-50 (type B1.1 & B1.3) and TIA/EIA-492CAAB fiber

**These SMF products are ideal for use with ISIS 10G links.**

### 1.8.3 Fibre Optic Transceivers for 10 Gigabit Ethernet

Avid supplies SFP+ transceiver modules for use in ISIS7000 V2.x engines (ISS2000) Both short range SR modules for used with multi-mode fibers and long range LR modules for used with single-mode fibres are available.

Avid supplied Foundry XFP transceiver modules for use in ISIS7000 V1.x engines (ISS1000). The same modules are available in an Avid supplied FESX 424 switch.

Avid Supplies Cisco X2 modules for used with Avid supplied Cisco Catalyst 4948-10GE devices. Both short range SR modules for used with multi-mode fibers and long range LR modules for used with single-mode fibres are available. The later generation C4948E uses SFP+ modules.

When using a customer supplied 6500 switch, Xenpak Modules with the X6704 10 Gigabit Ethernet interface card should be used. Both short range SR modules for used with multi-mode fibers and long range LR modules for used with single-mode fibres are supported.



### 1.8.4 Media Converters for Gigabit Ethernet

In some situations Zone 1 or Zone 2/3 Avid clients may be more than 90m from the Gigabit Ethernet switch and need to be individually connected by fiber optic connection. While the Foundry FESX-424 has 4 x SFP slots, this is not available on the Cisco C4948-10GE switch or the ISS blade. The newer C4948E has 4 x 1G/10G slots which can be populated by SFP/SFP+ as appropriate.

Some customers do not like to use media converters because they are unmanageable devices and each one needs its own PSU device.

This can be somewhat mitigated by a solution that offers rack mounting facilities and a common PSU for centralized deployment, and using a fiber optic NIC such as the Intel PRO1000MF or PRO1000PF.

Unlike Gigabit Ethernet switches, media converters do not need to be tested for their buffering capabilities, generally there either work or they fail. Some testing has been done on a suitable solution, which while not supplied by Avid, can be purchased by integrators for inclusion in the overall solution. The suggested product is multi-voltage, also there is good sales/support in Europe, North America, Latin America and Asia Pacific regions.

Details can be found on the Avid Knowledge base at the URL below:

<http://avid.custkb.com/avid/app/selfservice/search.jsp?DocId=244039>

### 1.8.5 Patching for Copper Structured Cabling

In some situations Zone 1, Zone 2 or Zone 3 Avid clients and servers may be directly patched to the local switch, but in other situations a formal patching system may be preferable.

There are so many variables to consider, and there is no clear reason why Zone 1, 2 and 3 connections should not go via patch panels, the deciding customer preference or logistical issues.

In a normal structured cabling system you have physical hops, NIC – WALL OUTLET (up to 5m), WALL OUTLET – PATCH (up to 90m), PATCH – NIC (up to 5m).

Providing these distance are not exceeded and the same correct category of cabling is used and all “legs” are tested and comply with relevant specifications there should be no detrimental effect for Ethernet signals.

A badly deployed system with substandard patches and too many physical hops can be a nightmare, but so can poor custom made direct cables, or badly routed cables.

Patching allows more flexibility for adds, moves and changes, however these element in Zone 1, 2 and 3 should be fairly static. However, with patching when a cable fails and the switches are in a different bay of racks, the patching is a massive advantage vs. dedicated cables.

The decision also depends on the types of switches. The needs of a single/pair of 49xx switches providing local ports is very different to a pair of fully stacked 6500 with 7 x 48 port Gigabit Ethernet and 2 x 8 port 10Gigabit Ethernet blades.

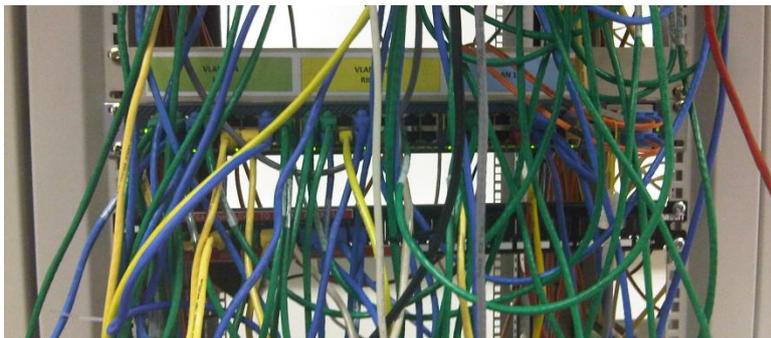


In a standard networking deployment, using Catalyst 6500 or Nexus 7000 or a stack of Catalyst 3750E, patch panels and numbered cables are highly recommend, probably essential! An Avid deployment should be no different in this respect.

For a small system with for 2 x 49xx in adjacent racks patch panels and numbered cables are not essential if done well, but still might be recommended.



A well laid out 6500 example



A 4948 example from less controlled environment. The 4948 is in there.... somewhere



## 1.9 IP Requirements

ISIS 7000 (Classic) systems requires a large IP space as each Engine requires 34 IP addresses, 17 in each VLAN , one for the switch blades and one each for the sixteen ISIS Server blades.

Although this section uses address from RFC 1918 192.168.0.0/16 private IP range (see [http://en.wikipedia.org/wiki/Private\\_network](http://en.wikipedia.org/wiki/Private_network)), any valid IP range is permitted. IP ranges do not have to be contiguous although this is advisable to reduce administrative complexity.

The IP requirements for ISIS5000 (V1.0 Release Q3/2010), are much reduced with a less than 10 IP addresses required all engines and embedded System Director.

### 1.9.1 Ranges required

The default addressing used is 192.168.10.0/24 and 192.168.20.0/24. This is used in all documentation. Ideally these ranged should not be used.

IP address use

- A fully loaded 8 Engine ISIS system needs 136 IP addresses in each VLAN, Plus addresses for System Directors and Media Manager
- For a system with 12 Engines this rise to 204 addresses per VLAN for 12 Engines R1.x
- This is required for Each ISIS VLAN
- This does not include others servers
- This does not include Interplay Production servers and Media Indexers
- This does nor include clients

Within the ranges IP addresses should be used in groups to keep similar devices in the same sub ranges.

A mid sized system with 6 engines will general use approx 200 IP addresses in each VLAN for Zone 1 and Zone 2 clients only. So this means two /24 networks are required as an ABSOLUTE MINIMUM. One or more additional ranges will be required for Zone 3 instances so this is why Avid requests a CIDR range of 1024 addresses.

A large system with 12 engines including supplementary servers and clients will use about 300 ip addresses in each VLAN for Zone 1 and Zone 2 clients only. So this means two /23 networks are required as an ABSOLUTE MINIMUM. One or more additional ranges will be required for Zone 3 instances so this is why Avid requests a CIDR range of 2048 addresses

Suggested use of Range

- For systems with up to 8 Engines Avid requests a CIDR block of 1024 addresses
  - Split into 4 x 256 Address blocks

- One For ISIS VLAN left
- One For ISIS VLAN Right
- Two blocks for Zone 3 VLANS as required such As interplay, Archive, Mezzanie
  - IF possible allocate a CIDR range of 2048 to cover future expansion possibilities.
- For systems with of 10 or 12 Engines Avid requests a CIDR block of 2048 addresses
  - This will be split in to one range of 512 addresses for each ISIS VLAN
  - The remaining addresses will be used for Zone 3 instances as required
- CIDR Blocks of contiguous addresses are preferred as enable summarized addressing which reduces routing table size
  - Non contiguous ranges are also acceptable but this means more work is required on the routing between the environments
- Very small systems (i.e. 2 engines max) can be done effectively in a block of 512 addresses divide in to 4 block of 128
- Other network sizes can be used but allocating two /16 networks is wasteful
  - This method will allow the 3<sup>rd</sup> octet of the address to be used to indicate device type



Note: DUAL STACK ISIS SYSTEMS of more than 12 Engines.

Beginning ISIS V2.4 systems of up to 20 Engines are permitted. Each ISIS stack or management domain, requires its own contiguous address space in a common subnet for each VLAN, but the two “ranges” do not have to be contiguous to each other. When building a dual stack system you must ensure that there is sufficient gap for each stack to become 12 engines. Each ISIS VLAN can still be achieved in a /23 subnet leaving approx 80 IP addresses for supporting servers, but it encourage the use of Zone 3 connected editors unless a /22 range is allocated per ISIS VLAN.

### 1.9.2 Default IP Ranges

Using the documentation default IP ranges (192.168.10.0/24, 192.168.20.0/24, 192.168.30.0/24) is not recommended as this is not a CIDR (Classless Inter Domain Routing) (<http://en.wikipedia.org/wiki/CIDR>) range. While 10/20/30 will work satisfactorily for a standalone system, should the system be integrated with a house network this will increase the administrative overhead required? Using a CIDR range reduces the size and administration/configuration associated with router tables and interconnecting networks.

Note deliberate used of differing RFC 1918 Private IP ranges in following examples.

If the customer does not provide a range, then for a small system that will never grow beyond 8 engines, use 192.168.32.0/22 this will provide the following:

IP NETWORK	VLAN NAME	VLAN NUMBER	COMMENT
------------	-----------	-------------	---------



NUMBER			
192.168.32.0/24	ISIS LEFT	VLAN32	MASK 255.255.255.0
192.168.33.0/24	ISIS RIGHT	VLAN33	
192.168.34.0/24	INTERPLAY	VLAN34	
192.168.35.0/24	ZONE 3B/OTHER	VLAN35	

For a medium system which will grow up to 12 engines, use 172.18.32.0/21 this will provide the following:

IP NETWORK NUMBER	VLAN NAME	VLAN NUMBER	COMMENT
172.18.32.0/23  MASK 255.255.254.0	ISIS LEFT	VLAN32	Lower 256 block: ISIS and System Directors only Upper 256 block: Servers & Clients
172.18.34.0/23  MASK 255.255.254.0	ISIS RIGHT	VLAN34	Lower 256 block: ISIS and System Directors only Upper 256 block: Servers & Clients
172.18.36.0/24	INTERPLAY	VLAN36	
172.18.37.0/24	ZONE 3B/OTHER	VLAN37	
172.18.38.0/24	ARCHIVE	VLAN38	
172.18.39.0/24	FUTURE	VLAN39	

For a large dual stack system which will grow up to 20 engines, use 172.18.32.0/20 this will provide the following:

IP NETWORK NUMBER	VLAN NAME	VLAN NUMBER	COMMENT
172.18.32.0/22  MASK 255.255.252.0	ISIS LEFT	VLAN32	Block 1 <sup>(256)</sup> Servers Block 2 <sup>(256)</sup> ISIS STACK 1 Block 3 <sup>(256)</sup> ISIS STACK 2 Block 4 <sup>(256)</sup> Editors
172.18.36.0/22  MASK 255.255.252.0	ISIS RIGHT	VLAN36	Block 1 <sup>(256)</sup> Servers Block 2 <sup>(256)</sup> ISIS STACK 1 Block 3 <sup>(256)</sup> ISIS STACK 2 Block 4 <sup>(256)</sup> Editors
172.18.40.0/24	INTERPLAY	VLAN40	
172.18.41.0/24	ZONE 3B/OTHER	VLAN41	
172.18.42.0/24	ARCHIVE	VLAN42	
172.18.43.0/24	FUTURE	VLAN43	
172.18.44.0/22	FUTURE		

If a block of 16 addresses is not acceptable then 12 networks could be represented with two CIR ranges as:

172.18.32.0/21 and 172.18.40.0/22

### 1.9.3 VLAN numbering

Often the 3rd Octet will be used as part of the VLAN number, but this can create numbering challenges when variable sub-netting is deployed to conserve/optimize IP space.

IP NETWORK NUMBER	VLAN NAME	VLAN NUMBER	COMMENT
10.128.32.0/24	ISIS LEFT	VLAN32	MASK 255.255.255.0
10.128.33.0/24	ISIS RIGHT	VLAN33	
10.128.34.0/25	INTERPLAY	VLAN341	MASK 255.255.255.128
10.128.34.128/25	INEWS	VLAN342	
10.128.35.0/25	ZONE 3A/OTHER	VLAN351	
10.128.35.128/25	ZONE 3B/OTHER	VLAN352	

When the basic /24 subnet is broken down into smaller blocks, the simple concept of using the 3rd octet is broken by its own simplicity. As any one of the octets can never be more than 256, this means one method might be to use value above 300, somehow related to appropriate octet used by associated VLANs.

For simplicity, use with VLAN numbers 2-999. While other values are possible, there are vendor dependencies.

Also make sure to NAME the VLANs and use the same/similar descriptions on the virtual interfaces.

### 1.9.4 Routed Interconnecting Networks

For connecting between the Avid environment and corporate environment, routed links are required. These links are generally point-to-point /30 links with two active addresses. Depending on the configuration more than one link may be required.

```
interface GigabitEthernet1/48
  description UPLINK to CORP.
  no switchport
  ip address 172.16.100.1 255.255.255.252
```

### 1.9.5 Interconnecting Networks example

The example below builds upon a dual uplink example from a single device which also used an Etherchannel for added capacity of resilience. It also uses L3 routed interfaces with /30 ranges.

```
!
interface Port-channel20
  description PRIMARY UPLINK TO CORP
  no switchport
  ip address 172.16.100.1 255.255.255.252
!
```



```
interface Port-channel21
  description SECONDARY UPLINK TO CORP
  no switchport
  ip address 176.16.100.5 255.255.255.252
!
<OUTPUT OMITTED>
!
interface GigabitEthernet1/45
  description PRIMARY ETHERCHANNEL 1/2 to CORP.
  no switchport
  no ip address
  channel-group 20 mode on
!
interface GigabitEthernet1/46
  description PRIMARY ETHERCHANNEL 2/2 to CORP.
  no switchport
  no ip address
  channel-group 20 mode on
!
interface GigabitEthernet1/47
  description SECONDARY ETHERCHANNEL 1/2 to CORP.
  no switchport
  no ip address
  channel-group 21 mode on
!
interface GigabitEthernet1/48
  description SECONDARY ETHERCHANNEL 2/2 to CORP.
  no switchport
  no ip address
  channel-group 21 mode on
```

### 1.9.6 Using Static Routes and HSRP

As Avid sells Cisco C49xx products which have only the standard IP software image, hence cannot participate in the scalable routing protocol such as EIGRP or OSPF. This is not as inconvenient as it initially appears. As the Avid is normally an island, then all that is needed is a default route up toward the house network combined with a static route in the corporate/house switches pointing down to Avid island. Of course the static route must be redistributed into the corporate routing protocols. The diagram below shows the concept which applies equally to C4900M and C4948, plus shows how the different HSRP groups must be used at each end. And provide example configuration.







WS-C4948-10GE-S Cisco Catalyst 4948-10GE, Standard Multilayer Image (SMI) (RIP, static routes, IPX, AppleTalk).

Cisco CAT4xxx IOS ENTERPRISE SERVICES provide the full the routing capabilities supported by this platform. Precise details will vary with type and level/version of Cisco IOS so the appropriate data sheets and release notes should be checked for exact specifications.

Using a HSRP instance on the Avid switches with a default route and an HSRP instance on the House Switches with a static route is a simple method to avoid the need for routing protocols.

### **1.9.8 ISIS 5000 IP address use**

Unlike the ISIS 7000 the ISIS5000 is much more conservative in its use of IP addresses. In a multi-engine system each Engine has an address in a single subnet. At product launch V3.0, the system is limited to 40 licenses, hence an entire system, all the servers and clients, would fit into a single /26 subnet of 62 addresses. At version 3.2 this up to 90 clients and 6 engines were permissible which will require a /25 subnet. However it might be that future versions will have more licenses so to cope with future expansion possibilities a single /24 range of 254 addresses is recommended.

### **1.10 MAN/WAN Connections**

Client connectivity to Avid ISIS 7000 was designed as a LAN application. Therefore using an ISIS client across a MAN or WAN connection poses many technical challenges. It is achievable with the correct configuration and end-to-end path considerations, such as QoS.

Note: WAN (Wide Area Network) is a generic term and used to indicate connections between sites, but generally means a lower speed connections using the legacy technologies up to 45Mbps. The term MAN (Metropolitan Area Network) is generally applied to the faster ATM/SDH/SONET and Ethernet based technologies with speeds above 100Mbps, which now extend beyond the metropolitan boundaries.

Avid recommends a Gigabit Ethernet connection for editing clients, and this requirement applies to a MAN scenario too. Many providers can supply high quality Gigabit Ethernet MAN solutions. Providing latency and jitter are within Avid guidelines (re-stated below), such solutions are supportable. The number of clients which can be supported in a Gigabit Ethernet link depends on the video resolution, quantity of concurrent streams and the percentage of guaranteed bandwidth available (as this might be a shared link used for corporate Data and voice services). Also the burst profile of the editing client must be considered, and best practice network design guidelines define that a MAN/WAN link should not be designed to be constantly loaded above 50% capacity, when it reaches this level an upgrade should be planned and deployed before 75% average capacity is reached.



Avid Interplay Assist clients using MPEG-2 browse media can work at Fast Ethernet, but there are restrictions on client count. Providing latency and jitter are within Avid guidelines, such solutions are supportable but are **not recommended**. However, Gigabit Ethernet is a much more viable solution and the price difference for most providers between a Fast Ethernet and a Gigabit Ethernet link presentation is not ten-fold, and often fractional access such as Gigabit Ethernet presentation with a guarantee of 25% or 50% data throughput is much better than Fast Ethernet.

The table below shows that **5mS** is the maximum latency which should be considered acceptable for a Gigabit Ethernet connection; the impact of the same latency on a Fast Ethernet connection is more significant and the figures should be halved.

Value	Behavior	Comments
0ms	System performs on test network as if locally attached	
5ms	Noticeable degradation in scrubbing performance, slight delay in play function (minimal)	RECOMMENDED Maximum Jitter and Latency - combined
10ms	Particularly noticeable delay in scrubbing, 1s delay from pressing play to material playing, may not be suitable for editors	USEABLE
20ms	More noticeable delay in scrubbing, 2.5s delay from pressing play to material playing – this would most likely be unsuitable for editors	UNSUITABLE

Note that ISIS traffic is not suitable for inspection by firewall. This is explained elsewhere in this document.

Any solution requiring this workflow is subject to approval from Avid network consultants.

### Network Specifications needs for WG to WG

This is workflow dependant, based on resolution and multiplication factor of real time transfer requires between WG. End to end latency is a key factor in achievable performance. With a Fast Ethernet link or a Gigabit Ethernet link which has 5mS of latency they will only get a about 10MB/S (2 x IMX30), Which about tops out a Fast Ethernet link but leaves 90% unused capacity on a Gigabit Ethernet link, however it is possible to use RFC1323 scaling to overcome the limitations of the standard TCP algorithm on a Gigabit Ethernet link with latency.

Avid can offer assistance in optimizing link utilization on paid-for consultancy basis.

### Network Specifications for Remote Clients

The specifications are same as for a LAN client, Avid requires a Gigabit Ethernet connection with less that 9mS of latency, preferably less than 5mS. Then consider number of concurrent clients/streams. Fast Ethernet is not recommended, even for Low Res clients. While it can be made to work, latency has much more impact, and then there is buffering, which is very dependant on which switches will be used etc.! Other solutions such as IP KVM may be a better approach where BW is limited.



As for text based clients such as iNews, then E1 2Mbps MEGASTREAM is fine, but Interplay Access probably needs more like Fast Ethernet to ensure a responsive application with head frames.

This type of requirement, where the workflow will dictate the speed and latency prerequisites, while the available speed and latency will dictate the workflow capability. Generally this requires network consultancy.

### **1.10.1 MAN - Example deployment**

On a LAN deployment it is possible to support up to 12 DV25 streams in a single Gigabit Ethernet path proving the upstream Gigabit Ethernet switch has sufficient buffering and the ISIS client is set for low resolution. This would be halved to 6 streams if using the ISIS client set for medium resolution or if using DV50/IMX 50 media. If the MAN link is shared then this should be reduced further depending on the quantity of guaranteed bandwidth available. For example it may be possible to only support 4 x DV25 streams or 2 x IMX 50 streams concurrently.

The above example is based on Based on legacy 256KB chunk size, With ISIS V2.0 and the default chunk size of 512KB the quantity of clients supportable is reduce by 50%.

Any solution requiring this workflow is subject to approval from Avid network consultants.

### **1.10.2 MAN - Proven deployment**

One successful deployment supported two NewsCutter editors at a remote site, connecting to and ISIS separated by approx 160Km (100miles) with a MAN link which traverses both land and sea. This link provided 1Gbps (initially it was 622Mbps) and had an average latency between of 5-6 ms on the primary path.

The solution as used for reviewing, editing and even occasionally ingesting material at DV25 until a second system was built at the remote site.

There were many challenges addressed by the customer IT team in terms of preventing this traffic from being subjected to the existing QoS and shaping profiles plus avoiding it being sent via a firewall. This site used MPLS and the devices at both ends were placed in the same VRF domain.

ISIS editors use UDP, but this link (actually the applications) was also tuned for file transfers using TCP and it was possible to consume almost 2/3 of link capacity for a single transfer representing more than 10x Real time workgroup-to-workgroup transfer.

## **1.11 DHCP**

Avoid DHCP IN Zones 1 and 2

DHCP is only useful where is a lot of “movement” in a network.

- Stick with STATICALLY assigned address for TransferManagers, NewsCutter etc.
- AirSpeed does not support DHCP

DHCP can be used for PC client if required

DHCP for Avid client devices in Zone3 will be dependant on customers’ policies.



DHCP will usually be default for Avid client devices in Zone 4 – customer network

## 1.12 10G Link aggregation

In large scale deployments, with a high number of Zone 4 video clients (also applies to Zone 2 and Zone 3) it may be necessary to aggregate 10G links from each ISIS 7000 VLAN toward the Border switch (or switches). ISIS does support the aggregation of 10G links using an Etherchannel style connection; however the full facilities of Etherchannel are not supported in V1.x software.

ISIS 1.x has a very BASIC implementation of Link Aggregation, it is capacity centric only. Basic Link aggregation adds capacity by making a single logical link with multiple physical links, if one physical path fails the whole logical link is affected. In ISIS 2.0 Avid supports the ability to work with a reduced number of links rather than just have a "leaking pipe" which renders the aggregate broken. In ISIS 1.x an aggregate link with 2 paths or 4 paths will be broken if 1 path fails. In ISIS 2.0 an aggregated path with 2 links will still run on 1 good path and one broken path (or 4 links will still run on 3 good paths and one broken path), hence a more complete algorithm in ISIS 2.x.

A combination of ISIS 1.x and Link aggregation and HSRP is a flawed design, because of the manual intervention is necessary at the ISIS to completely disable the broken link. Another approach to mitigating the design limitation would be to use GLBP (gateway load balancing protocol) instead of HSRP (hot standby routing protocol). This is a configuration feature on the Cisco switch and has NO COST IMPLICATIONS. While this has been "unofficially" tested by AVID with excellent results, it has not yet received official blessing.

The setup of a Cisco switch is described in this document referenced below. The set-up for one VLAN only is shown, and this would need to be replicated on the second ISIS VLAN

The Avid ISIS 7000 Ethernet Switch Reference Guide available at:  
[http://avid.force.com/pkb/articles/en\\_US/Compatibility/en348609](http://avid.force.com/pkb/articles/en_US/Compatibility/en348609)

### The settings in the Cisco 4948 or 6500 switch should be

```
interface Port-channel10
switchport
switchport access vlan 110
switchport mode access
.
.
.
interface TenGigabitEthernet1/49
switchport access vlan 110
switchport mode access
channel-group 10 mode on
!
interface TenGigabitEthernet1/50
switchport access vlan 110
switchport mode access
channel-group 10 mode on
.....
```





Similar commands in 6500 but with different interface numbers

The aggregate links has also been proven on a Foundry FESX 424 as a Zone 2 connection and the Cisco Catalyst 4948.

The foundry commands for aggregation in an FESX are

```
trunk ethe 25 to 26
trunk deploy
```

In a Foundry Super X the 10G interface would use Slot numbers, for example

```
trunk ethe 4/1 to 4/2
```

Link aggregation has also been tested successfully on the Foundry Big IRON RX, SX/SuperX and MLXe (v5.x s/w) across 10G I/O modules.

In ISIS 1.x it is necessary to keep link aggregation groups and an even number, i.e. 2, 4 or 8. An aggregate link with 8 channels is unlikely, and it the maximum permitted by Etherchannel. In ISIS 2.x Avid supports a more complete link aggregation algorithm with ability to work with a reduced number of links or even configured with an odd number for normal operation.

### **1.13 Deploying Transfer Manager**

In ISIS 1.x Transfer Manager (applied equally to Interplay Transfer and Archive Provider) has always required a Zone 1 connection to both ISIS VLANs. A third NIC with a different IP network is required to exchange data with the destination system other wise the bandwidth gets transmitted across the ISIS backplane twice! This can cause some design challenges and the binding order of NICs need to be controlled in combination with registry based persistent routes so the 3<sup>rd</sup> NIC is the preferred exit path to appropriate foreign networks, and it is best if the 3<sup>rd</sup> NIC is on a different PCI bus to those connecting with ISIS.

When a 3<sup>rd</sup> NIC is not viable, this “double booking” can be overcome by connecting the Transfer Manager as Zone 2 connections. This is not officially supported by Avid Engineering at this time but many sites do this with out any problems, as long as the bandwidth planning per ISS is done correctly it works the same. This means that external BW does not need a 3<sup>rd</sup> NIC and does not travel twice over the ISIS backplane and the egress point will now be the Zone 2/3 switch. It does however add extra load to the NIC used to exchange traffic with ISIS and will also add to the PCI bus loading.

Also see information in Section 2.8 for further information on triple NIC deployment.

### **1.14 Jumbo Frames and legacy applications**

ISIS 7000 does not support jumbo frames. While ISIS 7000 might be a more efficient system if it used Jumbo frames, to exchange data with clients, instead it uses a highly fragment datagrams based on 1500 byte MTU. There must be a reason. Jumbo frames are a great tool in a controlled and restricted environment, such as a server room where all devices are running on a small number of switches. But the administrative overheads of supporting jumbo frames across a wider network are significant. Unlike regular IP datagrams which can

be fragment by a router in the end-to-end path and re-assembled by the end host, jumbo frames cannot be fragmented by a layer 2 device (and some Layer 3 devices) in the same way, when a jumbo frame encounters a non jumbo frame enable layer 2 interface it will be dropped, hence a jumbo frame travelling toward a ISIS will have a curtailed journey. Successful jumbo frames don't just happen, both end hosts have to be configured for jumbo frames on the NIC and both end applications have to be configured for jumbo frames, and the capability is negotiated at the beginning of an IP exchange.

The only Avid application that is likely to use jumbo frames is AIRSPACE playout and ingest server. Used in a Media network environment it needed an intermediate Transfer Manager which was connected with Fibre Channel on one side and Gigabit Ethernet on the other, normally by an Alacritech Gigabit Ethernet (ToE) NIC then onto the AIRSPACE, The connection between the AIRSPACE and the Transfer manager may have been a directly cross over cable or a layer 2 switch which was Jumbo frame enabled in between. A very controlled and restricted environment!

In some cases the legacy AIRSPACE will be connected with ISIS. Ideally as described in section 1.12 above, this will be via a third network interface, suitably configured. But AIRSPACE can communicate without using jumbo frames. It is not as efficient and may affect performance but it still works.

### **1.15 Avid Low Res Encoder**

This device is the only product in the ISIS/Interplay Production family which MUST HAVE a Fast Ethernet connection. It is only supported on the Cisco Catalyst 4948, 4900M and the Foundry FESX 424.

It should work just as well on a:

- Cisco Catalyst 4500C with SUP 5 10G and a WS-X4506 Interface card,
- Cisco Catalyst 4500E with SUP 6 and a WS-X4648 Interface card
- Foundry FESX 448 ( with appropriate QD settings)
- Foundry Super X ( with appropriate QD settings)

The Low Res Encoder is known to fail on when connected via Cisco Catalyst 6500 using the only approved Gigabit Ethernet interface card WS-X6748. This is due to insufficient egress buffering on this card for a Fast Ethernet connected Avid video client.

### **1.16 Resilient First Hop Protocol**

There are several methods for providing a fault-tolerant default gateway. The Avid switch configuration examples describe the use of HSRP (Hot Standby Router Protocol) for Cisco solutions and VRRP (Virtual Router Redundancy Protocol) for Brocade/Foundry solutions, requiring the use of two layer 3 switches. When a third VLAN exists on both switches, this will not be connected to ISIS, so a switch to switch interlink must exist to carry this VLAN. Ideally there should be a switch to switch interlink trunk carrying all VLANS, with ISIS LEFT and ISIS RIGHT VLANS having increased path cost on this sw-sw link so it is normally blocked. This permits communication between the HSRP/VRRP instances in the event of a 10G path (to ISIS) failure. This switch to switch interlink should also be "protected" via an Etherchannel or secondary path (normally blocked by spanning tree) as it is of critical importance when supporting a clustered Interplay Production deployment.



This switch to switch interlink is already addressed in 4900M Config E using a 20G Etherchannel but does not feature in Config E for Cisco 4948 or Brocade FESX 424. Similar principles to the 4900M example can be used on the 4948 (and FESX424) to create an aggregate link between the switches.

NOTE: Cisco offers GLBP (Gateway Load Balancing Protocol) as an improved alternative to HSRP. While this has not been formally tested by Avid, it has been deployed successfully by several customers. One major benefit of GLBP vs. HSRP is that the improved protocol provides load-sharing to ensure both (all) the active interfaces are used to carry traffic. Another apparent advantage is a smoother transition in the event of a group-member failure. Wikipedia has some compact article which may be helpful.

### **1.17 INTERSWITCH link for resilient configurations.**

When using Config E with 4900M the standard config specified an inter-switch link across a 10G path. The primary purpose of this is to protect the Zone 2 clients in the event of a local 10G link failure to ISIS, and the ISIS VLANs have an increase path cost on these ports to block them (with STP/RSTP) under normal circumstances, but it also serves as a backup path for resilient first hop protocol (HSRP/GLBP/VRRP) intercommunications in event of a primary path (via ISIS) failure.

Note the 4948 config E does not have this feature, but it can be added using 1 G links (or aggregated depending on potential load) and when using Interplay Cluster a link for that VLAN must exist between the two switches.

The ISIS Config E is of course primarily about ISIS, but this will normally mean Interplay Production too, most likely in a cluster configuration, so it is important to include these VLANs on the inter-switch links as well. Any VLAN which exists with a resilient first hop protocol (HSRP/GLBP/VRRP) must be permitted on the inter-switch link.

Ideally this inter-switch link needs to be backup path for the ISIS VLANs too, so one method to do this is to use 2 x 10G aggregate and this is more important when using Nx10G aggregates from ISIS. But sometimes 2 x 10G might be excessive, so it might be preferable to back up the 10G link with a 1G link which should normally get blocked by STP (would need an increase port cost to ensure the desired path is normally blocked by STP/RSTP), so normally active is the 10G path passing non ISIS VLAN(s) traffic, if a 10G link to ISIS fails then the 10G inter-switch link will support the additional traffic (although some ISS 10G port oversubscription may occur depending on system planning), if the 10G inter-switch link fails but all other 10G links (to ISIS) are fine then the 1G link becomes active, to support the lower data rate transactional traffic.

### **1.18 RSTP settings for FHRP implementations**

The diagram below shows the basic principle for balancing VLANs across two switches using a First Hop Resiliency Protocol (FHRP) at layer 3 and Rapid Spanning Tree Protocol (RSTP) at layer 2. The VLANs are split into two groups, A & B. The RSTP root primary and root secondary parameters for those VLANs are associated with specific switches, to ensure a predictable convergence of the RSTP, these are also matched by the FHRP priorities for each VLAN, to ensure predictable operation of Active/Master and Standby states and functions.

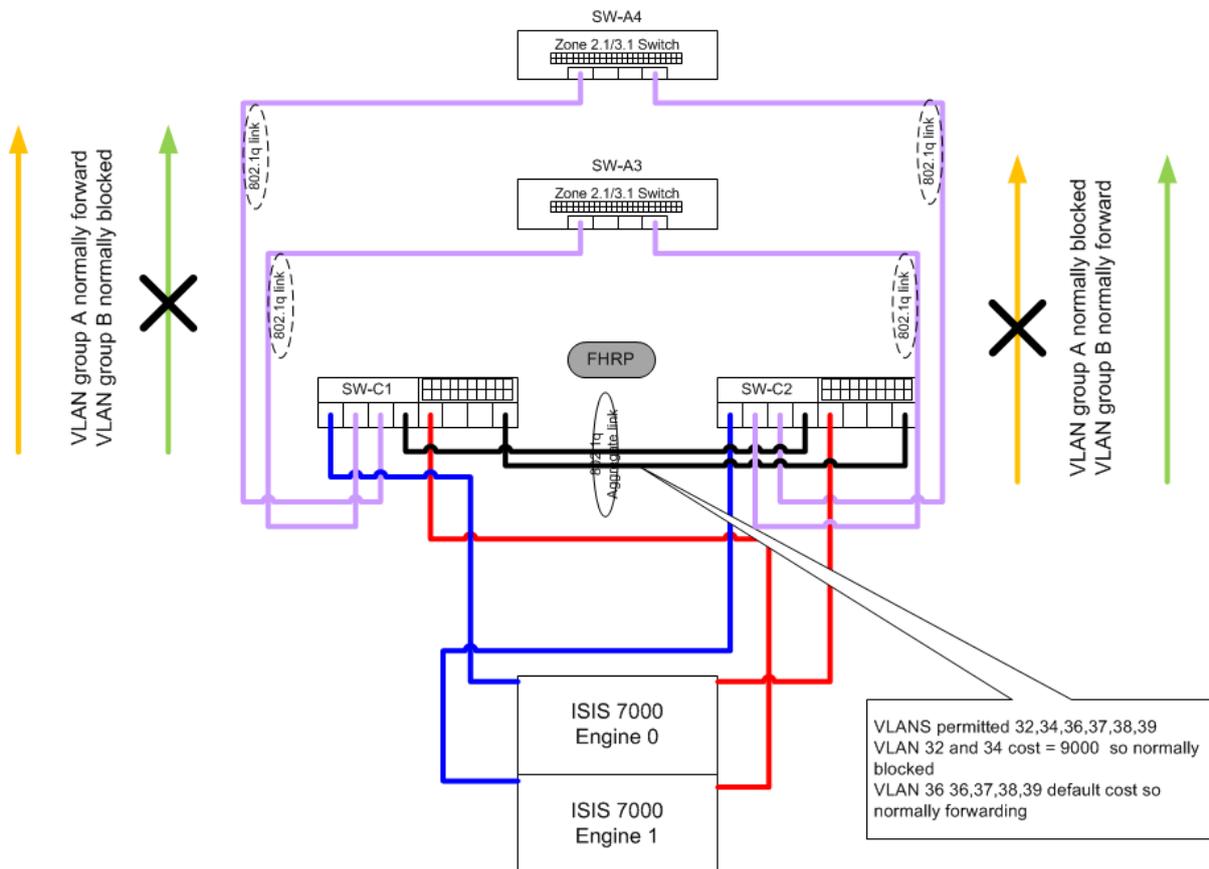
The example used RTSP with long costing.

The First Hop Resiliency Protocol (FHRP) may be:

1. Standards based Virtual Router Redundancy Protocol (VRRP)
2. Hot Standby Router Protocol (HSRP) - Cisco's initial, proprietary standard,
3. Gateway Load Balancing Protocol (GLBP) - a more recent proprietary standard from Cisco that permits load balancing as well as redundancy.

Under normal operations the group A VLANs will be forwarded from SW1 while blocked from SW2, and the group B VLANs will be forwarded from SW2 while blocked from SW1. Should the link between SW-C1 and SW-A4 fail then all VLANS would be forwarded from SWC-2.

This ensures that all paths are carrying traffic and that if over-subscribed 10G interface ports are used in the core switch that the optimum use is normally operational and that oversubscription will only be a possibility during the combination of heavy load and a path fault condition.



VLAN GROUP A RSTP PRIMARY SW1 RSTP SECONDARY SW2	VLAN 32	ISIS LEFT	VLAN GROUP A FHRP= ACTIVE SW1 Standby SW2
	VLAN 34	ISIS RIGHT	
	VLAN 36	Interplay & iNews	
VLAN GROUP B RSTP PRIMARY SW2 RSTP SECONDARY SW1	VLAN 37	Archive	VLAN GROUP B FHRP= ACTIVE SW2 Standby SW1
	VLAN 38	Editor	
	VLAN 39	Journalist	



The ISIS VLANs will be normally blocked on the SW-SW interlink unless one of the paths toward ISIS fails, in which case that VLAN will then pass across the SW-SW interlink, to support local Layer 2 and layer 3 requirements. The network design should be sized so that the ALL the links between ISIS and one core switch can fail, and the remaining paths from ISIS can carry the full load.

This section does not describe the advance interface/object tracking technique that can be applied to make the FHRP active status change based on link/path failures

## 2.0 Interplay Production Requirements

Interplay Production is a complex system with many requirements, which are variable depending on the system design. Below a subset of those requirements are discussed, but one suggested document is the Interplay Best Practices guide.

At the time of writing the most up-to date version is V2.6 available from

Avid Interplay Best Practices v2.6 [June 22, 2012]

[http://avid.force.com/pkb/articles/en\\_US/User\\_Guide/Avid-Interplay-Best-Practices-v2-6?retURL=%2Farticles%2Fen\\_US%2FUser\\_Guide%2Fen411843&popup=true](http://avid.force.com/pkb/articles/en_US/User_Guide/Avid-Interplay-Best-Practices-v2-6?retURL=%2Farticles%2Fen_US%2FUser_Guide%2Fen411843&popup=true)

Avid Interplay Best Practices v2.5 [November 18, 2011]

[http://avid.force.com/pkb/articles/en\\_US/User\\_Guide/en424651](http://avid.force.com/pkb/articles/en_US/User_Guide/en424651)

Avid Interplay Best Practices v2.4 [June 30, 2011]

[http://avid.force.com/pkb/articles/en\\_US/User\\_Guide/en411843](http://avid.force.com/pkb/articles/en_US/User_Guide/en411843)

### 2.1 To Multicast or not to Multicast

If all devices are configured with unicast lookup addresses, then the multicast command are not required.

Most implementations used UNICAST (DEC 2010) addresses for all lookup services, as this speeds up the Avid Service Framework. The JINI architecture is still central to Interplay Production.

The multicast commands for layer 3 switches in the sections below are not longer commonly deployed.

Avid Interplay Production makes use of the JINI distributed computing environment, which employs IP multicast. Multicast is required by the look up services. By default we configure the Cisco 4948 to use the default Cisco implementation of PIM SPARSE-DENSE-MODE.



All Interplay Production devices are listening on port 4160 receives the same packet if no router blocks it. Multicast packets are sent to port 4160 using the following IP addresses:

224.0.1.84 jini-announcement

224.0.1.85 jini-request

See <http://www.iana.org/assignments/multicast-addresses> for further multicast listings.

JINI is not an acronym. JINI is the name for a distributed computing environment that offers “network plug and play” effectiveness. It is the core underlying inter-communication technology behind the Interplay Production Framework services and Applications. It is a Java network technology developed initially by Sun Micro Systems®.

<http://www.sun.com/software/jini/overview/index.xml>

For Interplay Production clients in the customer network that also need to access the Avid Service Framework, i.e. more than just look up services, when Multicast routing is not permitted on the house network, two alternatives exist for remote (zone 4) clients:

- Multicast repeater
- Direct unicast IP address configuration on the client

Since 2009 the preferred deployment method has been to run Multicast routing only on the switches that connect directly with ISIS and support the Interplay Production elements, i.e. allow multicast in Zone 1, 2 and 3, but there is not requirement for multicast connectivity in Zone 4. Zone 4 clients are best configured with the IP address of the look-up server(s) (LUS), and not configured as part of the Avid Service Framework environment.

### 2.1.1 Multicast repeater - LEGACY

This is a (low specification) PC which accepts the multicast packets and then sends them via unicast IP to another multicast repeater in the Avid environment. When a house network is broken down into multiple VLANs and multicast routing is not repeated, each VLAN will need its own multicast repeater.

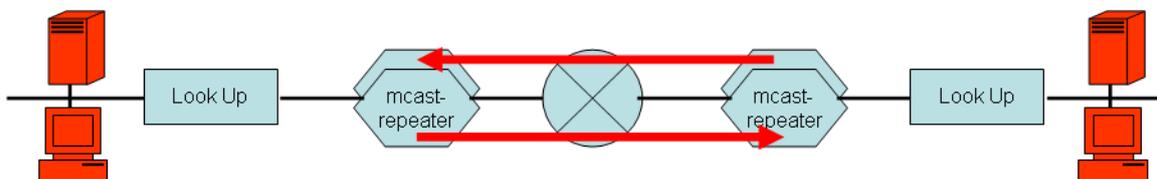


Figure 3 - Avid Multicast Repeater for Interplay Production

### 2.1.2 Direct Client Configuration

When a number of clients are widely dispersed, using a multicast repeater in each case is not feasible, individual clients can be configured with the unicast address of the look up service. As multiple IP address can be entered, this reduces some aspects of the in-built resilience of the JINI infrastructure.

Since 2009 this is now the preferred method for Zone 4 clients.

Ideally these distant clients will not participate in the Avid Service Framework, so have no requirement for Multicast features.





More details on Multicast requirements and recommendations can be found at:  
[http://avid.force.com/pkb/articles/en\\_US/Troubleshooting/en266811](http://avid.force.com/pkb/articles/en_US/Troubleshooting/en266811)

Avid Multicast Repeater Overview doc id: 266811

Description A Multicast repeater (MCR) is an Interplay Production software component used to bridge multicast packets that are required by the Interplay Production framework to send messages across IP subnets that would not otherwise pass multicast packets.

### 2.1.3 ALL Client Configuration Unicast

If all devices are configured with unicast lookup addresses, then the multicast commands are not required.

## 2.2 Multicast commands

The commands to enable multicast routing on the Cisco and Foundry switches are not included in the Standard ISIS configurations files, because these are for ISIS and not Interplay Production. However when ISIS and Interplay Production are used together it is recommended to enable the multicast routing commands below are an extract from the document “Avid Multicast Repeater Overview – Whitepaper”

[http://avid.force.com/pkb/articles/en\\_US/Troubleshooting/en266811](http://avid.force.com/pkb/articles/en_US/Troubleshooting/en266811)

Note that Multicast routing was never currently supported on the SMC8724 platform, which is shipped with V3.1.1.56 firmware. [2010 – SMC 8274 OBSELETE ]

### 2.2.1 Cisco Commands for Multicast

By default a layer 3 switch such as the 4948 used by Avid “contains” the multicast traffic in the local domains. This is a feature of all router boundaries and not specific to Cisco. To allow multicasts traffic to pass between networks, multicast routing must be enabled. An example of the configuration command needed on the Zone 2/3 border switch are given below. However multicast routing may also be required for Zone 4 clients in the customer network, unless an MCR is used. However, if the customer has multiple users located in many VLANS, as per a standard data network design, then several MCR instances may be required.

```
switch>enable
switch# show ip multicast
    Verify multicast status (off)
switch # config t
switch(config)# ip multicast-routing
!
switch config)# int vlan10
switch(config-if)# ip pim sparse-dense-mode
!
```



```
switch(config)# int vlan20
switch(config-if)# ip pim sparse-dense-mode
!
switch(config)# int vlan30
switch(config-if)# ip pim sparse-dense-mode
!
switch(config)# int vlan40
switch(config-if #) ip pim sparse-dense-mode
!
! Note - Repeat for other VLANs as required/configured (e.g. 30)
!
switch(config-if #)exit
!
switch(config)#exit
switch # show ip multicast
  Verify multicast status (on)
!
  switch # show run
  Verify Vlan setup shows pim setup
!
switch # copy run start
```

### 2.2.1 Foundry Cisco Commands for Multicast

By default a layer 3 switch such as the FESX-424 used by Avid “contains” the multicast traffic in the local domains. To allow multicasts traffic to pass between networks, multicast routing must be enabled. An example of the configuration command needed on the Zone 2/3 border switch are given below. However multicast routing may also be required for Zone 4 clients in the customer network unless an MCR is used. However, if the customer has multiple users located in many VLANS, as per a standard data network design, then several MCR instances may be required.

```
switch> enable
switch # config t
switch(config)# router pim
!
```



```
switch(config-if)# int ve 10
switch(config-if)# ip pim
!
switch(config-if)# int ve 20
switch(config-if)# ip pim
!
switch(config-if)# int ve 30
switch(config-if)# ip pim
!
switch(config-if)# int ve 40
switch(config-if)# ip pim
!
```

! Note - Repeat for other VLANs as required/configured (e.g. 30)

```
switch(config-if)# exit
!
switch(config)# exit
!
switch # write mem
```

### 2.2.3 ALL Client Configuration Unicast

If all devices are configured with unicast lookup addresses, then the multicast commands are not required. (This is a deliberate copy of section 2.1.3)

## 2.3 DNS

The availability of a local DNS instance is imperative for a successful Interplay Production deployment. Forward and reverse lookup is required. The DNS should contain zones for both ISIS VLANs and any other Avid-specific VLANs, such as transfer VLANs, Cluster VLANs, etc.

If a local DNS is not available, and DNS is provided via the corporate network it means that the Interplay PAM/ISIS environment cannot be disconnected from the corporate network in the event of a security issue. If the link to the corporate network is broken, either due to deliberate action or a failure, the ISIS and Interplay Production will not function correctly.

For optimum performance, Interplay Production/ISIS requires a DNS lookup with negligible latency; this is another reason why a local DNS is preferable.

In a new Interplay Production installation, a popular option is to have DNS installed and configured for all “Avid VLANs” on Avid General Purpose Servers. These servers can be configured a few different ways, but the most common are to either have each server dual-connected, one port to each primary ISIS VLAN, or to configure each server with a switch fault tolerant teamed connection, in a third, routed VLAN (to two different switches). Because the General Purpose Servers also run Avid Services Framework, and ASF relies





heavily on a single, primary connection, the teamed approach may be the best option for redundancy. These options should be discussed with an Avid Network Specialist at the time of planning and commissioning.

If integration with an existing corporate DNS structure is required, one viable option would be to setup DNS servers within the ISIS/Interplay Production core to serve these zones, and setup forwarders to forward all other unknown lookups to the corporate DNS servers.

### **2.3.1 Why is FQDN resolution required?**

Common understanding is that it's a 'security' feature of the 'auto-magic' Java discovery layer, used by JINI which is a component of Avid Interplay Production.

Fundamentally it is a requirement for the network protocol security that Avid use for AIF (Avid Interplay Production Framework). To get technical, the requirement has its basis in how Java implements Remote Method Invocation protocol (RMI).

Basically what happens under the covers is that when a remote client (like Interplay Assist) tries to contact a service (like Media Indexer), the connection goes through some security checks before it is accepted. One of the security checks is to validate that the caller really is who he says he is. The protocol handler does that by looking at the IP address of the caller and doing a reverse DNS lookup to resolve a name; that name is then compared with the name the caller provided as part of the initial request. If they match, the system trusts that the caller is really who he claims to be. If the names differ, the call is rejected (because the caller isn't who he claims to be). FQDNs are used throughout so that we always know that name comparisons match; use FQDNs everywhere is preferable so that short name vs. long name mismatches are avoided.

A simple or short name is considered 'ambiguous' whereas the FQDN is 'unambiguous' in that it points to the exact location in the DNS hierarchy from the top-level down to the nameserver. Also some DNS resolvers treat the 'trailing dot' in a FQDN or lack of, differently. The correct method is to use a trailing dot as part of a FQDN. Wikipedia provides some further references.

<http://en.wikipedia.org/wiki/FQDN>

## **2.4 Active Directory with Interplay Production Cluster**

When a Clustered solution is deployed it is essential to have Microsoft Active Directory services available. This is because Microsoft Cluster, which is the underlying service providing a resilient Interplay Production solution requires Microsoft Active Directory services.

More details can be found in the "Avid Interplay Engine Failover Guide, For Cluster installations" available from the URL directly below which the most up to date document at the time of writing.

Avid Interplay Engine Failover Guide for AS3000 Servers  
Updated January 10, 2012





[http://avid.force.com/pkb/articles/en\\_US/User\\_Guide/en418451](http://avid.force.com/pkb/articles/en_US/User_Guide/en418451)

Avid Interplay Engine Failover Guide for SR2400 and SR2500 Servers  
December 12, 2011

[http://avid.force.com/pkb/articles/en\\_US/User\\_Guide/en429075](http://avid.force.com/pkb/articles/en_US/User_Guide/en429075)

Avid Interplay Engine Failover Guide v2.3 [387015]

Avid Interplay Engine Failover Guide • 0130-07643-02 Rev I • December 2010 • Created 12/10/10

<http://avid.custkb.com/avid/app/selfservice/search.jsp?DocId=387015>

## 2.4.1 Avid Interplay Production Active Directory Considerations

With the advent of Interplay Production and particularly the Microsoft Cluster solutions for both Asset Manager and Archive Manager, an Active Directory (AD) account is required. This account is used on both cluster nodes and allows the systems to failover in the event of a hardware or software problem.

Most large customer sites will already have an Active Directory containing all users, computers and other peripheral devices. The problem with integrating into this AD is strict security policies are likely to exist to restrict user access and computer functionality. These policies can cause Avid editing applications to fail as they require high levels of access on the local computer and a relatively unrestricted network connection to shared storage and asset management systems. In addition to this IT specialists that maintain the AD are not best placed to maintain critical broadcast systems.

Another reason to locate an AD instance within the Avid environment is security of operation. If the connection to corporate network is broken for any reason, whether deliberately or as a result of a failure, then Interplay Production will be greatly affected because the cluster will malfunction so it is **STRONGLY RECOMMENDED** that an AD instance exists within the Avid environment.

All of the above has led Avid to suggest three standard configurations for use when Microsoft Clustering is specified in the system, these are.

1. Avid supplies one or two servers to act as DNS and Active Directory servers. The second server offers redundancy, an important consideration in critical broadcast systems. Within the AD three organizational units (OU) would be created (Clients, Users and Servers.) This would allow policies to be deployed to clients and users, for example restricting internet access, and would also ensure critical server systems never inherited these policies. With this configuration it would be preferable if all computers were domain members and all logins managed from the AD.
2. Avid supplies one or more servers to act as DNS and AD. The AD would be very basic with only the cluster service account being used. All server systems would reside in an Avid workgroup and local logins would be used. Clients would be added to the customer AD, allowing centralized management of users etc by the IT departments.



3. Avid provide no dedicated hardware and rely on the customer infrastructure for both DNS and AD. In this environment it's strongly advised a customer AD server is connected inside the Avid segment of the customer network. In addition to that the avid related computers and users should be segregated into separate OU's which have limited policies applied.

In all the above examples the system directors are recommended to be left outside the Active Directory.

### **Active Directory Trust Relationship**

A one way non transitive outgoing trust would be the recommendation when linking the Avid Domain to the customers existing domain,

By making the trust non transitive you have to explicitly setup between two separate domain forests.

By making the trust a one way outgoing trust from the Avid domain you allow users in the customer domain access to Avid resources but Avid user have no rights in the customer domain,

Below is extract from Active Directory Overview document from Microsoft.

Technical\_Overview\_of\_Windows\_Server\_2003\_Active\_Directory.doc

### **Active Directory Basics**

Active Directory is the directory service for Windows Standard Server, Windows Enterprise Server, and Windows Datacenter Server. (Active Directory cannot be run on Windows Web Server but it can manage any computer running Windows Web Server.) Active Directory stores information about objects on the network and makes this information easy for administrators and users to find and use. Active Directory uses a structured data store as the basis for a logical, hierarchical organization of directory information.

### **Directory Data Store**

This data store is often simply referred to as the directory. The directory contains information about objects such as users, groups, computers, domains, organizational units (OUs), and security policies. This information can be published for use by users and administrators. The directory is stored on servers known as domain controllers and can be accessed by network applications or services. A domain can have one or more domain controllers. Each domain controller has a writeable copy of the directory for the domain in which it is located. Changes made to the directory are replicated from the originating domain controller to other domain controllers in the domain, domain tree, or forest. Because the directory is replicated, and because each domain controller has a writeable copy of the directory, the directory is highly available to users and administrators throughout the domain. Directory data is stored in the Ntds.dit file on the domain controller. It is recommended that this file is stored on an NTFS partition. Some data is stored in the directory database file, and some data is stored in a replicated file system, like logon scripts and Group Policies.

A useful online (Microsoft) resource for technical articles is



<http://technet2.microsoft.com/windowsserver/en/library/6f8a7c80-45fc-4916-80d9-16e6d46241f91033.msp?mfr=true>

## **2.5 Time-code and NTP.**

In the avid world it's critical that the CaptureManager has the timecode card and is locked to house LTC. This is because the CaptureManager will have control of devices such as airspeeds and low res encoders, also with LTC and both must be exactly the same, to ensure accurate recordings. The CaptureManager has the TC card and its local clock is locked to that card. Typically the CaptureManager is then made into an NTP server for all non framework clients, i.e. Airspeeds, Countdown etc.

Any PC that has the Interplay Production Framework time service running on it **MUST** have the windows time service disabled so that the two do not interfere with each other. In fact what happens is that the Interplay Production Framework time sync will stop so as not to cause clock hopping. Any client that has Interplay Production applications loaded will be running framework services even when Interplay Production applications are not running, because the Framework exists as a service. Interplay Production Framework services update the time every 10 seconds which is more frequent than NTP defaults or Windows time defaults.

Typically the local Active Directory domain controller (only mandatory if running an Interplay Production cluster) will be locked to CaptureManager, either via the NTP service or if it's running an Avid service such as the lookup service then it can use the Interplay Production Framework time sync.

When a CaptureManager is delivered via Manufacturing (Dublin) it comes with a pre-installed Time code card from Adrienne Electronics, <http://www.adrielec.com/pci-tc.htm>, model **PCI-LTC/RDR** is supplied, and pre-set with specific drivers which must not be changed.

As part of System Setup, Avid engineers configure the NTP server on the CaptureManager and link it with the Interplay Production Framework Services.

The **Avid CaptureManager Installation and Configuration Guide, v4.0** available at

[http://avid.force.com/pkb/articles/en\\_US/User\\_Guide/en266859](http://avid.force.com/pkb/articles/en_US/User_Guide/en266859)

### **2.5.1 Time Synchronisation for Avid Interplay™ systems**

Another very useful document can be obtained from Avid Knowledge base:

[http://avid.force.com/pkb/articles/en\\_US/How\\_To/en369915](http://avid.force.com/pkb/articles/en_US/How_To/en369915)

#### Time Synchronisation for Avid Interplay™ systems

Accurate time synchronisation is critical for all Avid Interplay and workgroup systems. For reliable operation, all devices which create media (such as editors, AirSpeeds, TransferManagers, etc) or which control the creation of media (such as CaptureManager servers, CaptureManager clients, etc) must have their time synchronised to within 3 seconds of incoming house timecode and to each other. To





keep the time-stamp of the created media files consistent, the shared storage must also be synchronised.

## **2.6 Interplay Production Assist Browse resolution.**

Avid recommend a Gigabit Ethernet connection for all clients. When using MPEG II browse resolution it is possible to use Fast Ethernet, but such application need the agreement and approval of Avid network consultants. See sections 1.5 and 1.8 above.

## **2.7 DHCP**

Avoid DHCP IN Zones 1 and 2

DHCP is only useful where is a lot of “movement” in a network

- Stick with STATICALLY assigned address for Interplay Production Media indexers, Transcode servers.
- AirSpeed does not support DHCP

DHCP can be used for PC client if required, i.e. Interplay Assist or Interplay Access  
DHCP for Avid client devices in Zone3 will be dependent on customers’ policies.  
DHCP will usually be default for Avid client devices in Zone 4 – customer network

## **2.8 Streaming Server deployment practices**

There are many considerations on how to deploy the Streaming Server, such as:

- Should it be Zone 1 or Zone 2 connected.
- Should it have a different external path for the Access clients?
- What switches are need in the corporate network

### **2.8.1 Network Zones and DNS**

This is both a topology issue and a DNS issue. Since most corporations have integrated DNS, the remote and the local users are using the same DNS.

For the One Customer installation the Project Manager has taken the following approach ...

- Triple connected Stream & Delivery servers, dual connected to ISIS VLAN (10 & 20) and teamed connection to an ‘outbound’ VLAN (40).
- Set the ‘outbound’ as Primary.
- Only registered the ‘outbound’ in DNS.
- Used a few static route statements to tell the boxes to route to local VLANs via 10 & 20 connections.

This config requires that the Streaming & Delivery Server(s) be able to route from their 40 connections to find FRAMEWORK components, like LUS. It also means that all users, local and remote, will ‘find’ the outbound connection.

It also offers the following advantages ...





- Dedicated pipes for local ISIS connection and for remote streaming & delivery.
- Simplified traffic shaping, all routes to remote subnets will be made through the 'outbound' primary.
- Simplified DNS ... all users find the 'outbound' connection.

### **2.8.2 Network Requirements for Interplay Access streaming clients.**

The objective of the Interplay Streaming server for Interplay Access clients is to greatly reduce the network path and workstation requirements, and even to permit browsing via WAN grade connections. There are several servers involved in the solution which connect with ISIS, depending on the elements used in the solution there may be Stream Server, Publish Server and Delivery Server.

The Interplay Access clients receive an Ultra Low Bandwidth stream, which does not use fragmented packets like ISIS clients, nor any form of oversubscription. The Quick-Time wrapped MPEG 4 stream with Compressed Audio requires approximately 1Mbps. Both the uplink path from the Avid environment, and the edge device demands are very low.

Basically, Interplay Access clients receive a "web-class" stream which has been demonstrated over 802.11 WiFi, and can even be used via a moderate broadband/DSL connection. There should be no issues co-existing with VoIP deployments. Ideally in a LAN deployment for this type of client will be Fast Ethernet minimum.

Consider that a single stream is approx 1Mbps, so a 100 of them would consume a Fast Ethernet pipe, and take approx 10% capacity of a Gigabit Ethernet pipe. Of course depending on the network design for resilience, there may be more than a single Gigabit Ethernet interface between Avid and House network. Plus remember that this does not take into account any other traffic using the uplink path such as Transfer Engine or Delivery Server.

### **2.8.3 Firewall Parameters for Interplay Stream Server Clients**

The resources below describe the requirements for Interplay Stream Clients when using Interplay Stream Server introduced Interplay 2.0.

<http://www.soundscreen.com/streaming/firewall.html>

Open the appropriate ports on the Firewall. This allows the streaming server to be accessed via RTSP/RTP on the default ports and provides better use of network resources, lower speeds for client connections and less load on the server. The Ports that need to be open for unrestricted streaming include:

TCP Port 80: Used for signaling and streaming RTSP/HTTP (if enabled on server)

TCP Port 554: Used for RTSP

UDP Ports 6970 - 9999: used for UDP streaming.

Note a smaller range of UDP ports can usually be used (typically 6970-6999).

TCP Port 7070: Optionally used for RTSP (this port is used by Real Server, and QuickTime/Darwin can also be configured to use this port)

<http://www.geeklan.co.uk/?p=14>

## 2.8.4 Firewall Parameters for Interplay Streaming Server Clients

The resources below describe the requirements for Interplay Stream Clients when using Interplay Streaming Server introduced Interplay 2.4.

To enable streaming from the Interplay Streaming Server to Interplay Access clients, the network has to allow UDP and TCP traffic on the following ports:

- Direct client/server communications: 554, 7070 TCP
- UDP streaming: 6970 - 6999 UDP
- Interplay Engine communication: 80 TCP

This port usage applies to both SR2500 and AS3000 servers.



Note: If UDP packets are lost on the way from the server to the client there is no direct feedback. The Streaming Server does not get any feedback that the packets were lost, and the Access client must wait until a timeout occurs.

## 2.8.5 Supported Config

At the time of writing this section, the only tested/supported config is 100 clients with a single streaming server. So with this you only need 1 NIC for ISIS and 1 NIC for the clients.

The Streaming Server is based upon an open-source solution which does indicate the ability to scale with multiple servers for several hundreds of clients.

## 2.9 Interplay Production Copy Server

This section is supplementary to information in Avid Interplay Media Services Setup and User's Guide, and concentrates on the hardware deployment practices.

The COPY server is a unidirectional transfer service that requires all the normal Interplay Production elements to exist in both workgroups. If data transfer is required in both directions then two copy servers are required. The copy server has a Layer 2 relationship with ISIS. Layer 3 connections are not supported.

The COPY server is supplied (and normally configured) with 2 x 10G interface cards which would normally be connected as 10 G Zone 1 device, directly to one ISS on each ISIS system. It is possible to connect using 1 G interfaces but this will impact the performance achieved. The only testing which has been done is using Zone 1 connection at 10G.



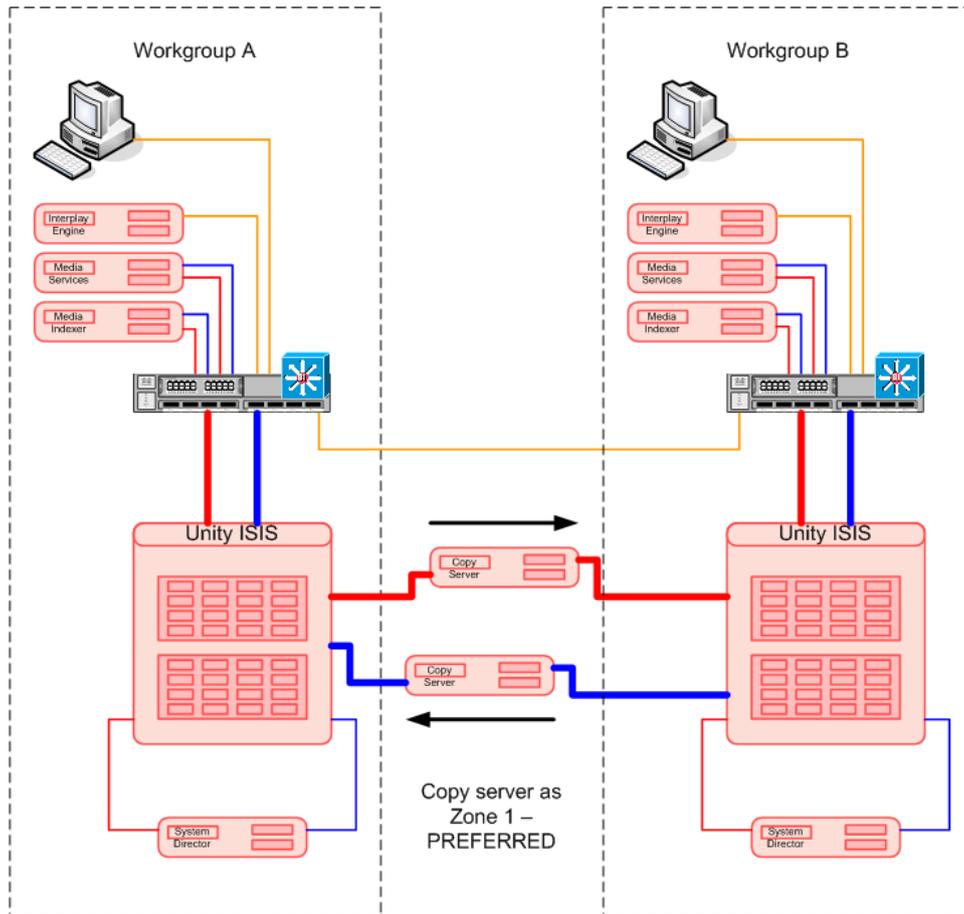


Figure 4 - Copy Server in Zone 1 - Preferred

Careful bandwidth planning is necessary to ensure an engine or ISS is not overloaded. Depending on the parameters associated with file transfer(s), the achieved throughput may vary from 30 to 300MB/S. Also consider that if necessary the bandwidth available to this device can be limited at the client level.



Note: One Engine should not be used to host 2 Copy Servers.

### Deploying in Zone 2

There may be some instances when the copy server cannot be connected as Zone1, but can be connected as Zone 2. This method has not been explicitly tested. The client itself has no understanding of Zone 1 or Zone 2; both are a layer 2 relationship with ISIS.



Deploying Devices in Zone 2 when using CHELSIO 10G cards in an SR2500 server has encountered some field problems, which were overcome by replacing with (standard Avid issue) CHELSIO 10G interface card with a (standard Avid issue) MYRICOM 10G interface card, which is also supplied in the AS3000 server. Deployment of MYRICOM in SR2500 is not qualified (Q3/2011), but is due for testing Q4/2011



INFORMATION



A Copy Server with MRYICOM 0G card has been successfully deployed in Zone 3 with a single connection in Q3/2011. This is a supported configuration beginning Interplay 2.5

The example depicted below shows a two engine and a one engine ISIS. On both Workgroups A and B, all the available 10G Zone 1 links have already been used so 10G ports on the 4900M switch are used.

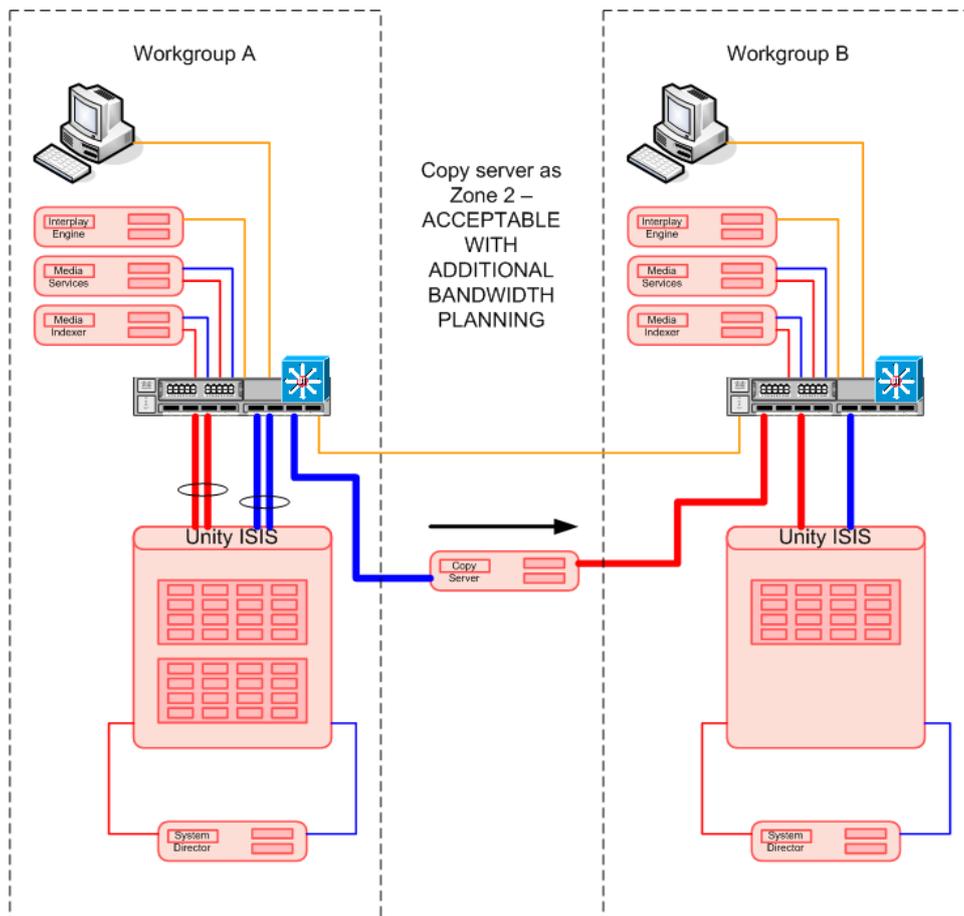


Figure 5 - Copy Server in Zone 2

In this situation careful bandwidth planning is critical to ensure an engine or ISS is not overloaded, however it may also be that the workflow dictate that COPY action do not occur during normal working hours when editors and ingest/payout are demanding bandwidth. Depending on the parameters associate with file transfer the achieved throughput may vary from 30 to 300MB/S. Also consider that if necessary the bandwidth available to this device can be managed at the client level.

A variation of this approach might be that Workgroup A has a 10 G Zone 1 connection and Workgroup B backup system is connected via 10G Zone 2. As stated above the clients are not aware of Zone 1 or Zone 2.



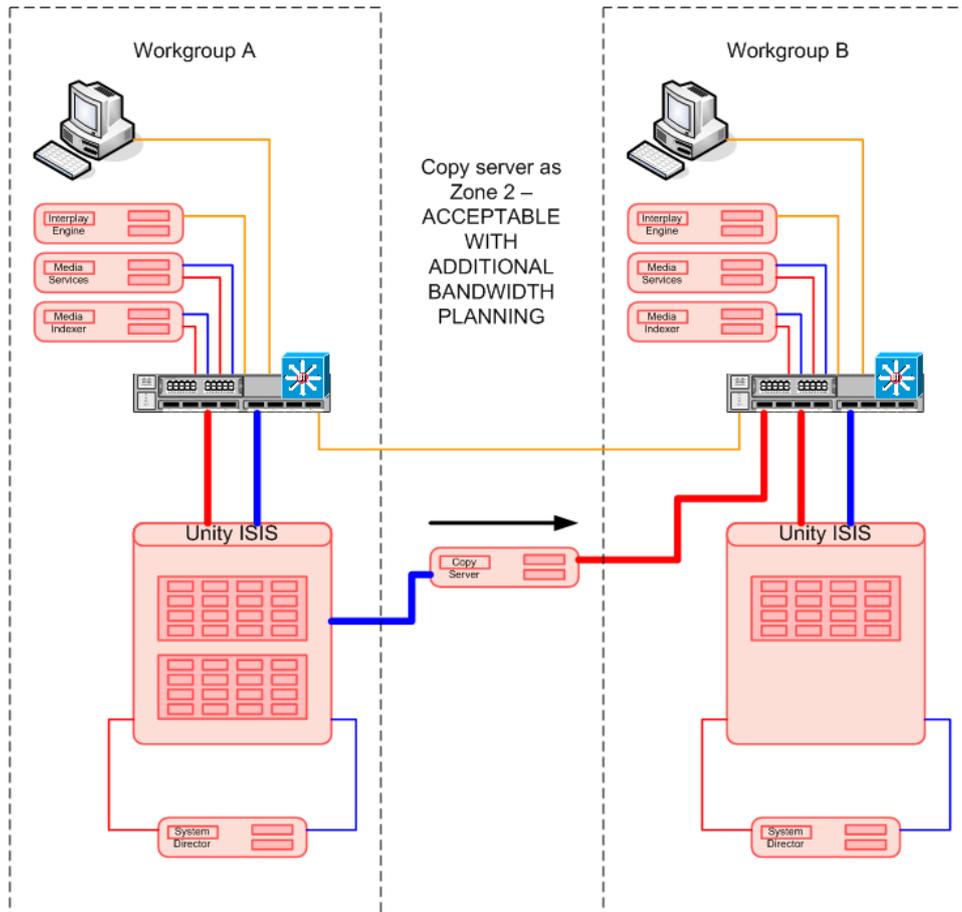


Figure 6 - Copy server in Zone 1 and Zone 2

Also possible but not shown is connecting one side as 10G and the other side at 1G, such a configuration would be most likely a temporary setup to migrate media & metadata between ISIS systems during a major upgrade procedure.

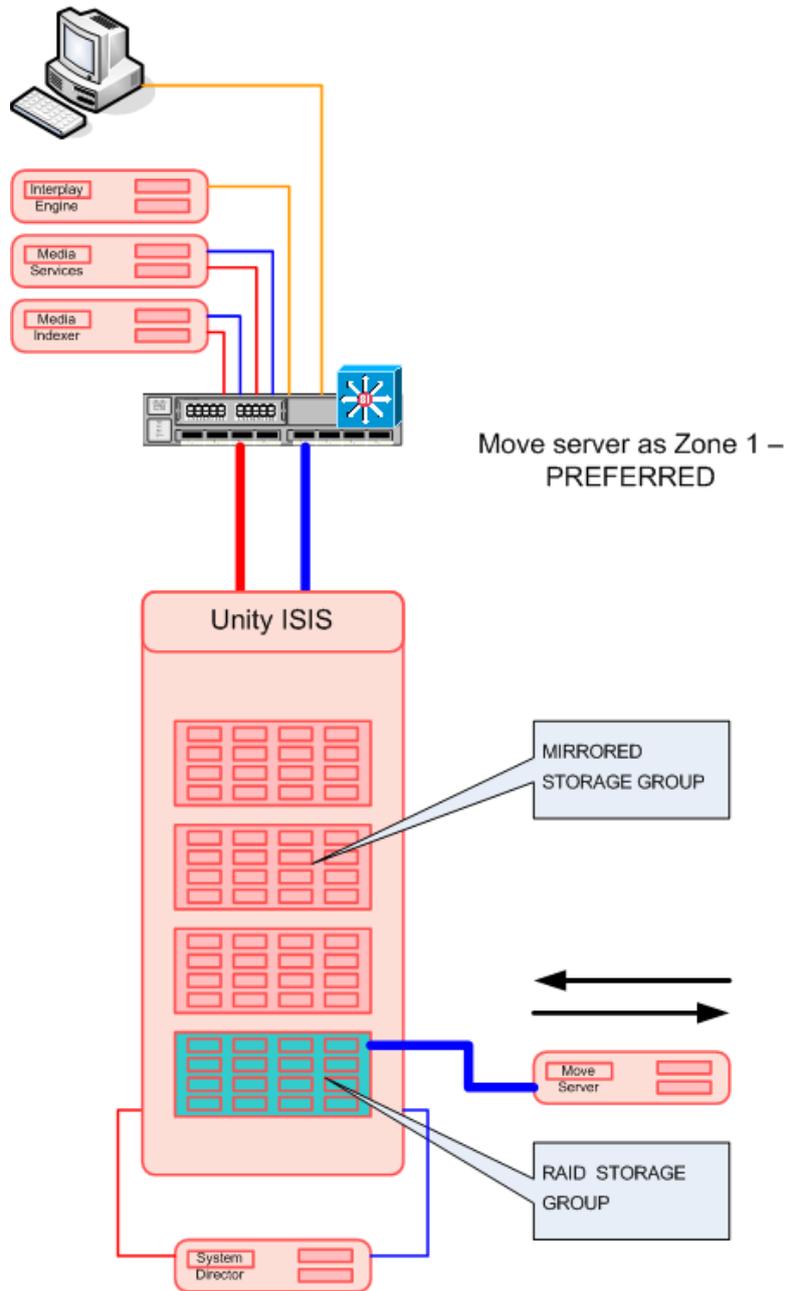
## 2.10 Interplay Production Move Server

This section is supplementary to information in Avid Interplay Media Services Setup and User's Guide, and concentrates on the hardware deployment practices.

The Move server can be used to move data between workspaces in the same storage group or between workspaces in different storage groups. This section will concentrate on the latter and considers a MIRRORRED storage group and a RAID storage group.

The MOVE server is supplied (and normally configured) with 1 x 10G interface card which would normally connect as 10 G Zone 1 device, directly to one ISS the ISIS system. It is possible to connect using 1 G interfaces but this will impact the performance achieved. The only testing which has been done is using Zone 1 connection at 10G.





**Figure 7 - MOVE server in Zone 1**

Careful bandwidth planning is necessary to ensure an engine or ISS is not overloaded. Depending on the parameters associated with file transfer(s), the achieved throughput may vary from 30 to 300MB/S. Also consider that if necessary the bandwidth available to this device can be limited at the client level.

When using a RAID Storage group and a MIRRORED storage group it is likely that the demand on the RAID storage group will be less, so this will be the best connection point for 10G interface card.

### Deploying in Zone 2



There may be some instances when the MOVE server cannot be connected as Zone1, but can be connected as Zone 2. This method has not been explicitly tested. The client itself has no understanding of Zone 1 or Zone 2; both are a layer 2 relationship with ISIS.

The example depicted below shows a three engine ISIS, where all the available 10G Zone 1 links have already been used so 10G ports on the 4900M switch are used for the Move server.

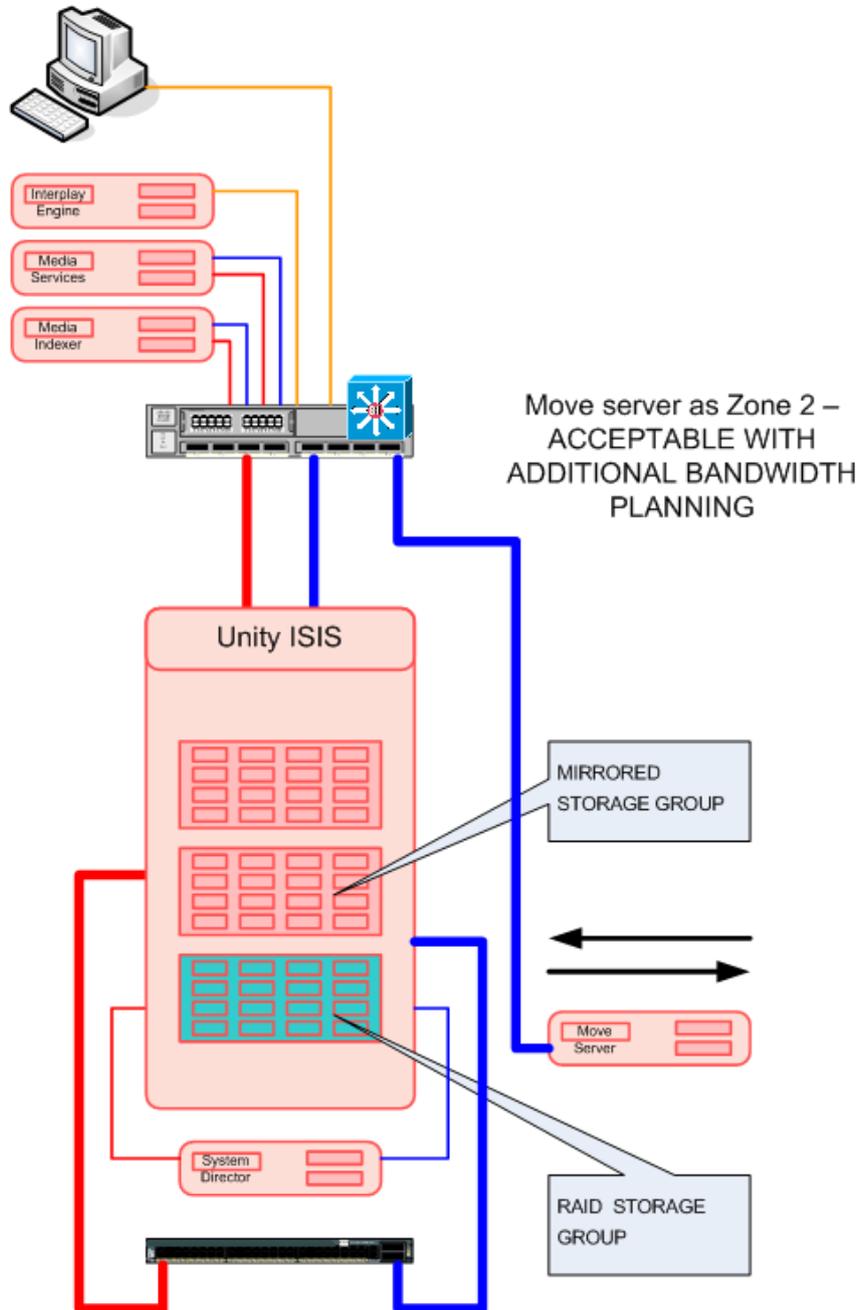


Figure 8 - Move Server in Zone 2

### Deploying as 1G

There may be some instances when Gigabit Ethernet interfaces are used, but this has not been tested or commonly deployed, but should work with restricted throughput.



Deploying Devices in Zone 2 when using CHELSIO 10G cards in an SR2500 server has encountered some field problems, which were overcome by replacing with (standard Avid issue) CHELSIO 10G interface card with a (standard Avid issue) MYRICOM 10G interface card, which is also supplied in the AS3000 server.

## 3.0 Enhancing Network Performance

Depending on what needs to be enhanced, different techniques may be deployed. By using the approved network switches (see section 1.3) and the correct network interface cards (or LoM chipsets) with recommended settings (see section 1.5) the network will already be significantly optimized for ISIS video traffic. There may be other elements of the overall solution which could be optimized such as wide area file transfers which may benefit for some of the information below.

### 3.1 TCP Window Sizing

Adjusting the TCP window size is a common method of improving network performance. This is where you set RWIN (RcvWindow), RWIN is the single most important tweak of TCP parameters. Raising RWIN from default: (8760 for Win95/98/98SE/NT and 17520 for WinME/2K/XP) can greatly improve download speeds. But without using scaling these changes are limited.

The highest RWIN you can use without Windows Scaling being turned on is 65535. Simply put, RFC 1323 Windows Scaling (TCP Extensions for High Performance, <http://tools.ietf.org/html/rfc1323>) is needed to enter any number higher than 65535. WinME/2K/XP does not need the vtcp.386 patch required by earlier MS Windows operating systems). Most users do not need to go higher than 65535. In MS Windows, Windows Scaling "Defaults" to off (same as No), however this setting may be of use for server to server connections.

ISIS video clients use a UDP protocol for exchanging video between the client and the servers, but the techniques described above and articles given below can be of great assistance with TCP based connections such as TransferManager data communication with external clients using high bandwidth links, as high capacity WAN links.

#### Description of Windows 2000 and Windows Server 2003 TCP Features

<http://support.microsoft.com/kb/224829>

#### Microsoft Technet: TCP Receive Window Auto-Tuning

<http://www.microsoft.com/technet/technetmag/issues/2007/01/CableGuy/default.aspx>

#### Enabling High Performance Data Transfers

<http://www.psc.edu/networking/projects/tcptune/>



### **TCP Tuning**

[http://en.wikipedia.org/wiki/TCP\\_Tuning](http://en.wikipedia.org/wiki/TCP_Tuning)

### **Dr TCP**

<http://www.dslreports.com/drtcp>

### **DRTCP: How do I use it and what are all these settings?**

<http://www.dslreports.com/faq/578>

## **3.2 Useful Knowledge Base articles**

### **3.2.1 Starbucks Fix for ISIS v1.0-1.4**

[http://avid.force.com/pkb/articles/en\\_US/Troubleshooting/en241963](http://avid.force.com/pkb/articles/en_US/Troubleshooting/en241963)

Note that from ISIS 1.5 the software set attends to the issue of older client software releases and the setup is not required when the client software is installed properly.

## **4.0 PC and MAC Requirements**

The general workstation requirements can be accessed via Avid website. The information is dynamic and changes with s/w versions and workstation availability for Hewlett Packard and Apple. Sometimes more comprehensive details will be found in the appropriate README file.

### **4.0.1 SEPTEMBER 2012 URLS**

At the time of revising this section (SEP 2012) the content below was available: Avid Qualified Systems and IO hardware for Media Composer 6.5, Symphony 6.5, Newscutter 10.5, Assist 2.5 - 2.7, and Instinct 4.0 - 4.1

[http://avid.force.com/pkb/articles/en\\_US/compatibility/en422411](http://avid.force.com/pkb/articles/en_US/compatibility/en422411)

### **4.0.2 AUGUST 2011 URLS**

At the time of revising this section (AUG 2011) the content below was available: Windows Specifications for Avid Media Composer v5.5.x, Avid NewsCutter v9.5.x, Avid Assist v2.3, and Avid Instinct v3.5 [362927]

[http://avid.force.com/pkb/articles/en\\_US/Compatibility/en362927](http://avid.force.com/pkb/articles/en_US/Compatibility/en362927)

Windows Specifications for Avid Media Composer v4.0, Avid NewsCutter v8.0, Avid Assist v2.5, and Avid Instinct v2.5



[http://avid.force.com/pkb/articles/en\\_US/Compatibility/en290223](http://avid.force.com/pkb/articles/en_US/Compatibility/en290223)

#### **4.1 Customer provided platforms**

Avid provides qualified workstations and specifications, but sometimes this does not match with the policies of the customer. Workstations have the benefit of a fixed configuration from the vendor, usually for 12-24 months, but this comes with increased cost. In many cases the cost of workstation cannot be justified for desktop clients running the s/w editing products. The key elements to ensure for a non-standard clients is the quality of the network interface card:

- (1) Choosing a PC with an Intel LoM implementation and there is a high chance of success, or fit an additional Intel Pro 1000 NIC, the single port card is quite inexpensive when ordered in bulk.
- (2) The quality of the Video card, fit an approved NVIDIA QUADRO card as described on the Avid website.
- (3) System bus architecture. Using a cheap motherboard with everything on the same bus, even if the video card and NIC are correct, would still have some serious performance issues, especially where I/O hardware is being used.

#### **4.2 Customer tested platforms - 2007**

This information is superseded by newer models but remains for reference many customers have deployed DC7800 and DC7900 successfully, inline with the details given above in section 4.1

The description for Interplay Assist and iNews Instinct does not describe specific products. Some customers have deployed these applications on the HP XW4x400 workstations, but workstations are very expensive compared to PCs.

Models successfully tested so far for use with Interplay Assist and iNews Instinct

HP Compaq DC7700 Mini Tower with NVIDIA Graphics Adapter

- Integrated Intel 82566DM Gigabit Network Connection
- NVIDIA Quadro NVS 285 (128MB DH) PCIe x16 VGA Card
- Integrated High Definition audio with Realtek 4-channel ALC262 codec
- Works with Assist, needs soundcard parameters changed from default

HP Compaq DX7300 Business PC with NVIDIA Graphics Adapter

- Integrated Intel 82566DM Gigabit Network Connection"
- NVIDIA Quadro NVS 285 (128MB DH) PCIe x16 VGA Card"
- Integrated High Definition audio with Realtek 4-channel ALC262 codec
- Works with Assist, needs soundcard parameters changed from default

Note: the Intel 82566DM Gigabit Network Connection is a later & similarly capable chipset to the Intel 82546 Gigabit Controller use in the Intel Pro 1000/MT interface card

The HP Compaq DC7800 is currently being considered by another customer/project and has the same NIC chipset as the DC 7700, however this used a later version of the NVIDIA graphics family and at the time of writing is still awaiting evaluation.

### **4.3 Imaging PC Clients**

To save time during installation, the use of disk imaging is often employed. Interplay Production system can be deployed from and image providing certain criteria are observed as described below.

#### **Drive Images can Cause Duplicate Service IDs**

Prior to creating drive images using software, such as Norton Ghost™ ensure that all Interplay Production Service IDs are removed to avoid conflicts with duplicate IDs when the images are restored to multiple computers.

Workaround: Before creating drive images, do the following:

1. Stop all Interplay Production services and applications using the instructions in the Avid Interplay Production Framework User's Guide.
2. Search the directory where Interplay Production software is installed and permanently delete any files with the extension **.serviceID**.
3. Ensure the drive image is created before any Interplay Production services or applications are run again.

The next time the services and applications run, they will automatically create new and unique Service IDs.

Note: Later Interplay Production software checks for duplicated service ID's and creates new ones based on host name so a rename of a host triggers a "corrective" action on these ID's. However, it is still preferable to delete before taking an image.

Note: Imaging clients such as Media Composer SOFT or NewsCutter SOFT must be done before the s/w license is applied, as the license is associated with multiple hardware elements.

### **4.4 ULTRA HIGH RESOLUTION Clients**

When Avid introduced Ultra High Resolution Clients in Q4/2009 the preferred deployment was Zone 1 only in an Avid Unity compatible single channel 10 GbE PCI-e network interface card (CHELSIO) with short range (SR) integrated optics. For use with xw8600 workstations and dual quad-core xw8400s.

In Q3/2010 introduced the Avid ISIS 7000 compatible single channel 10 GbE PCI-e network interface card (MYRICOM) with short range (SR) integrated optics. For use with HP Z400 or Z800 workstations, with Windows XP32, Windows Vista64, or Windows7.

Also for use with dual quad-core MacPro workstations running Snow Leopard.





Check Avid website for latest MAX O/S support on editing platforms.

## 5.0 Network Designs

The qualified network designs can be found at in the “Avid ISIS 7000 Ethernet Switch Reference Guide” available at:

Other network designs need to be qualified. Avid network consultants have experience of many successful network deployments, plus they have learned some valuable lessons from some deployments which encounters a few challenges.

Examples of suitable networks are given below.

## 5.0.1 Cisco 6500

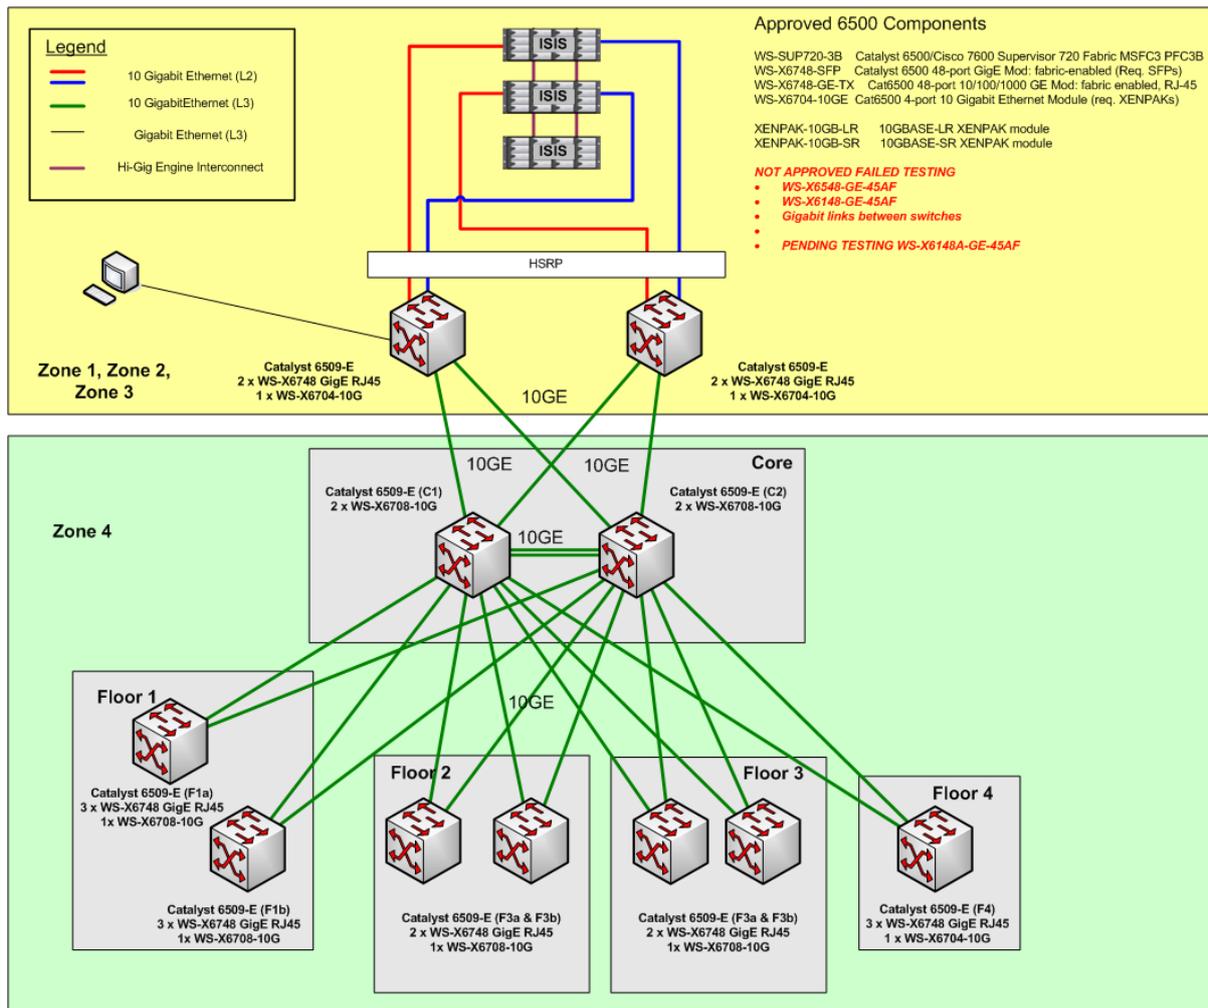


Figure 9 - Cisco 6500 Example

The diagram above shows a system with Dual Cisco 6500 switches working as an HSRP pair. The yellow section is the video production network and the green section is the corporate network. This example is based on the classical Hierarchical Design Model and concentrates on the Enterprise Campus sections of the Enterprise Composite Network Model. It does not consider the Edge distribution, Enterprise Edge, Service Provider Edge, Server Farms, or Management Network elements of the Enterprise Composite Network Model.

The 6500s which form the collapsed core and the 6500s which connect with ISIS will feature multiple PSUs and dual supervisor modules to ensure maximum availability. The edge switches may not have as much resilience built in depending on system requirements at this level. All interconnections between switches are 10G fibre optic to ensure non blocking characteristics and so that buffering is only required at the edge.

For larger systems the single 10G links from each ISIS VLAN may not be sufficient, using the 6500 means that aggregated links of 20G or 40G are possible if required.



Real Time Video is an aggressive traffic flow and not suitable to be subjected to the limitations of a firewall. By using a Cisco 6500 most of the PC based video clients can be placed in a “Zone 3” mezzanine network design which allows them to be located in the Video production network but still have access to the corporate infrastructure via an intermediate firewall.

### 5.0.2 Cisco 6500 and 4500

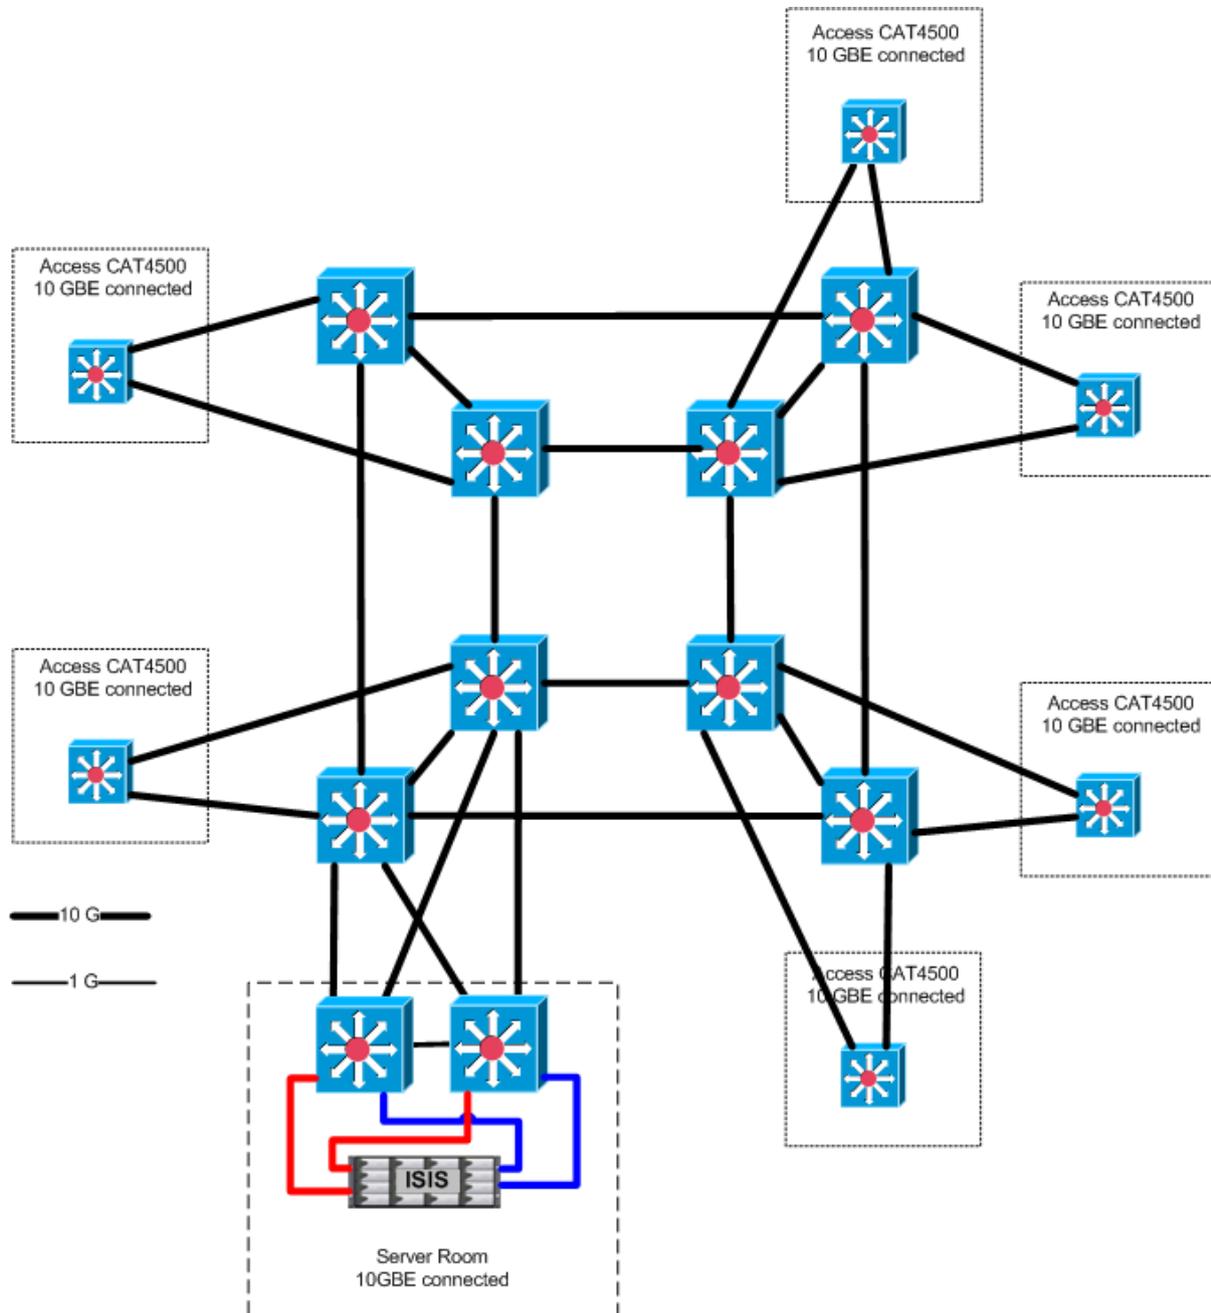


Figure 10 - Cisco 6500 and 4500 example

This network example uses a 6500 “double helix” core with 4500 at the access layer. End-to-end connectivity was exclusively 10 Gigabit Ethernet. ISIS connected with a dual HSRP pair of Catalyst 6500 using with Zone client Devices being Gigabit Ethernet connected on a WS-X6748 interface card. Dual stream editing clients in Zone 4 connected to an

uncontended WS-X4506 interface card with single stream browse clients devices connected to a contended WS-X4548 interface card, which also supported other corporate clients.

### 5.0.3 Cisco 4948 with cascaded 3750

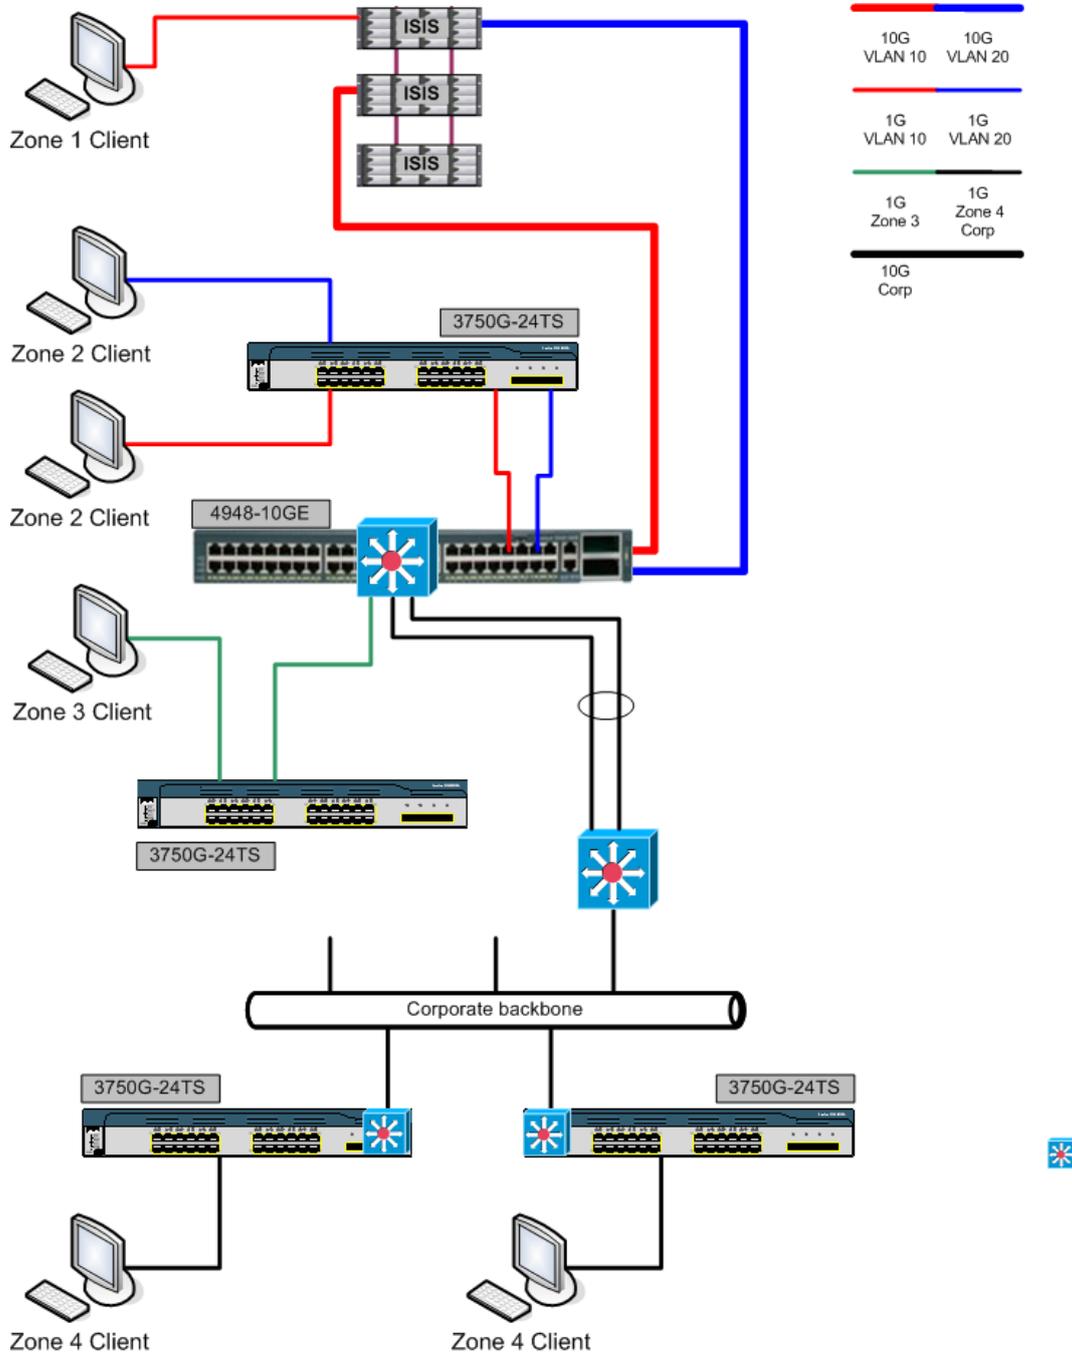
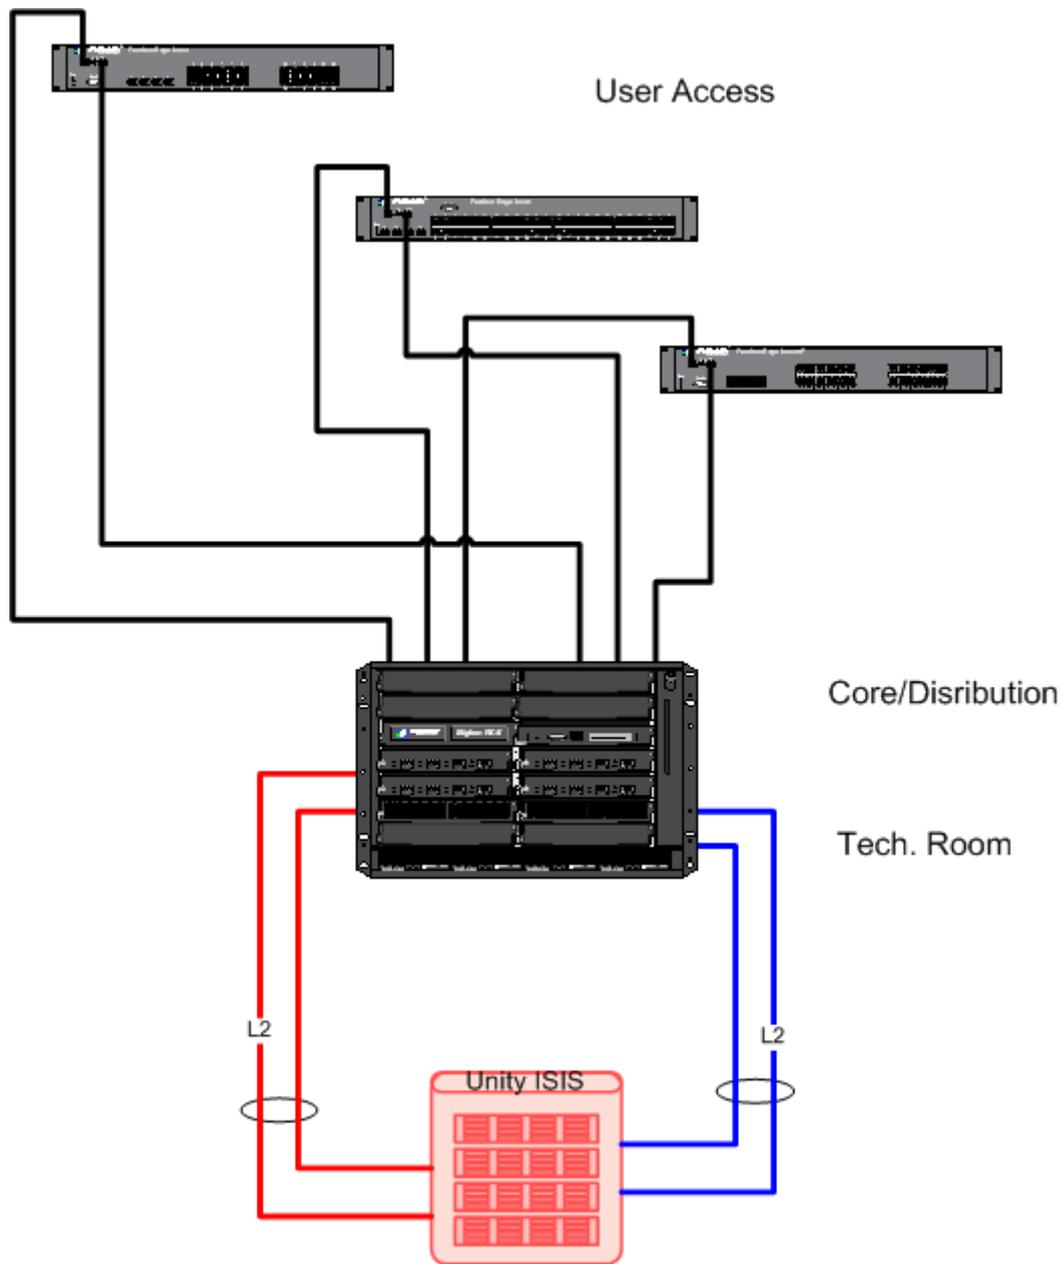


Figure 11 - Cascaded 3750G Example

This network example shows the deployment of Cisco Catalyst 3750 Gigabit Ethernet switches connecting multiple clients and relying on the buffering capability of the Catalyst 4948. This allows a RESTRICTED number of clients to work on the 3750, the quantity depends on the video resolution and the ISIS chunk size, it is not a fixed value so varies with every customer. One could use a 3560E-24 in place of the 3750E-24.



### 5.0.4 Foundry RX-8 core with FESX Edge

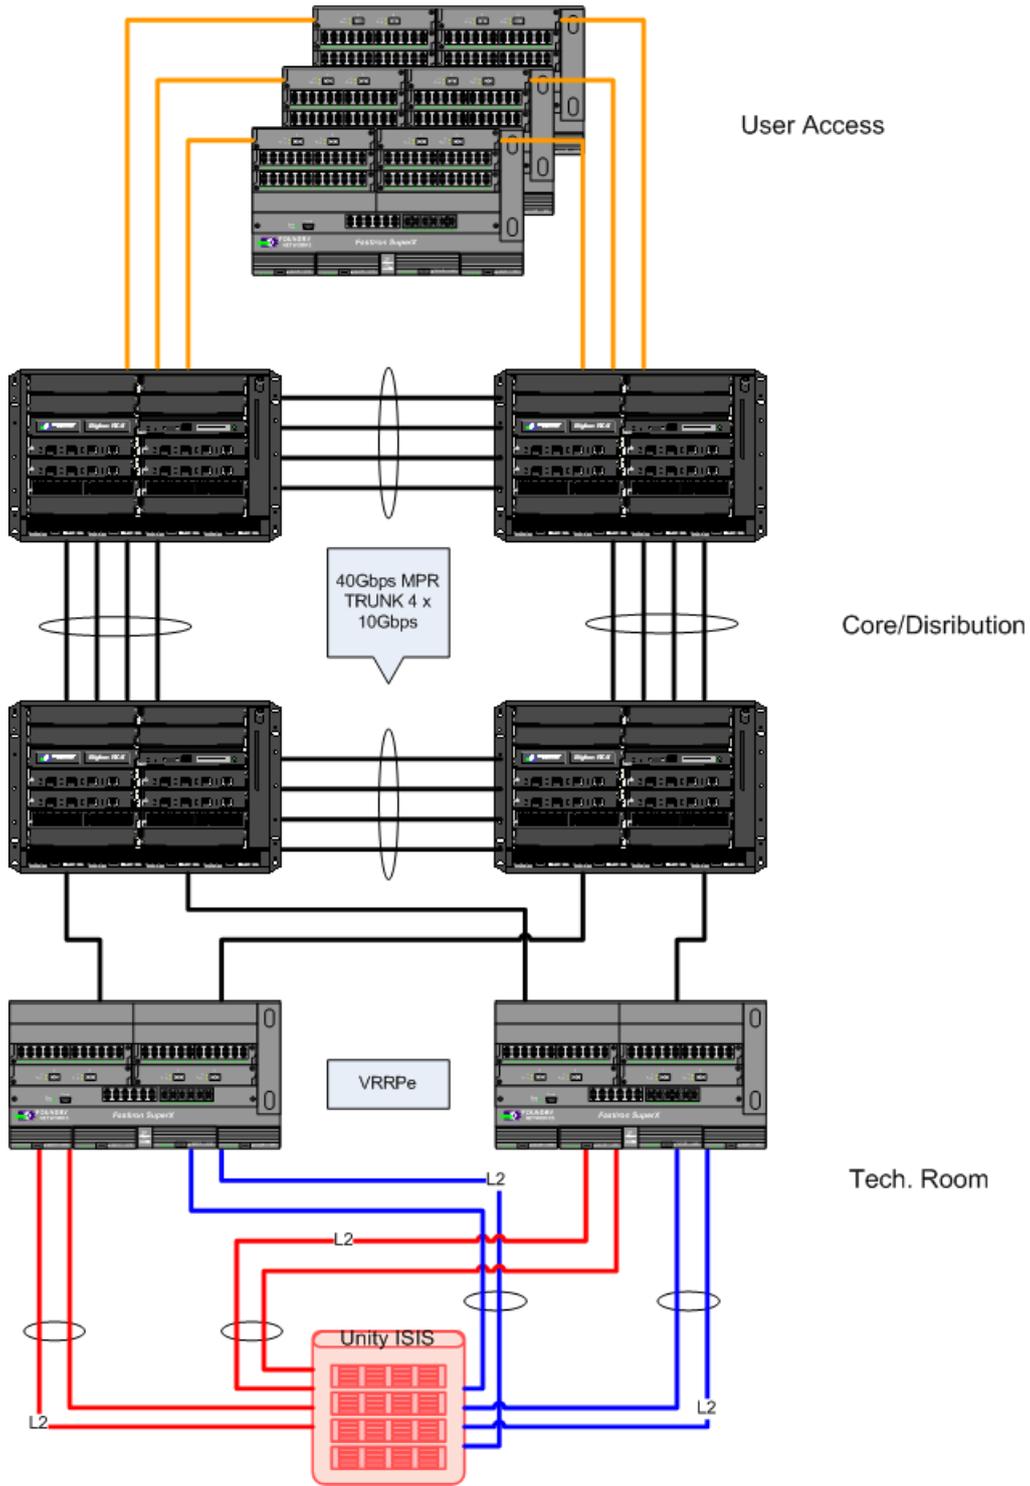


**Figure 12 – Foundry Example RX & FESX**

This network example uses a Big Iron RX switch as the heart of the system permitting easy expansion with link aggregation; the edge connectivity is fulfilled using FESX switches. In the diagram above the FESX 424 is a qualified product while the FESX 448 is an approved product. Also shown is a FESX 424HF which offer additional optical connections if required.



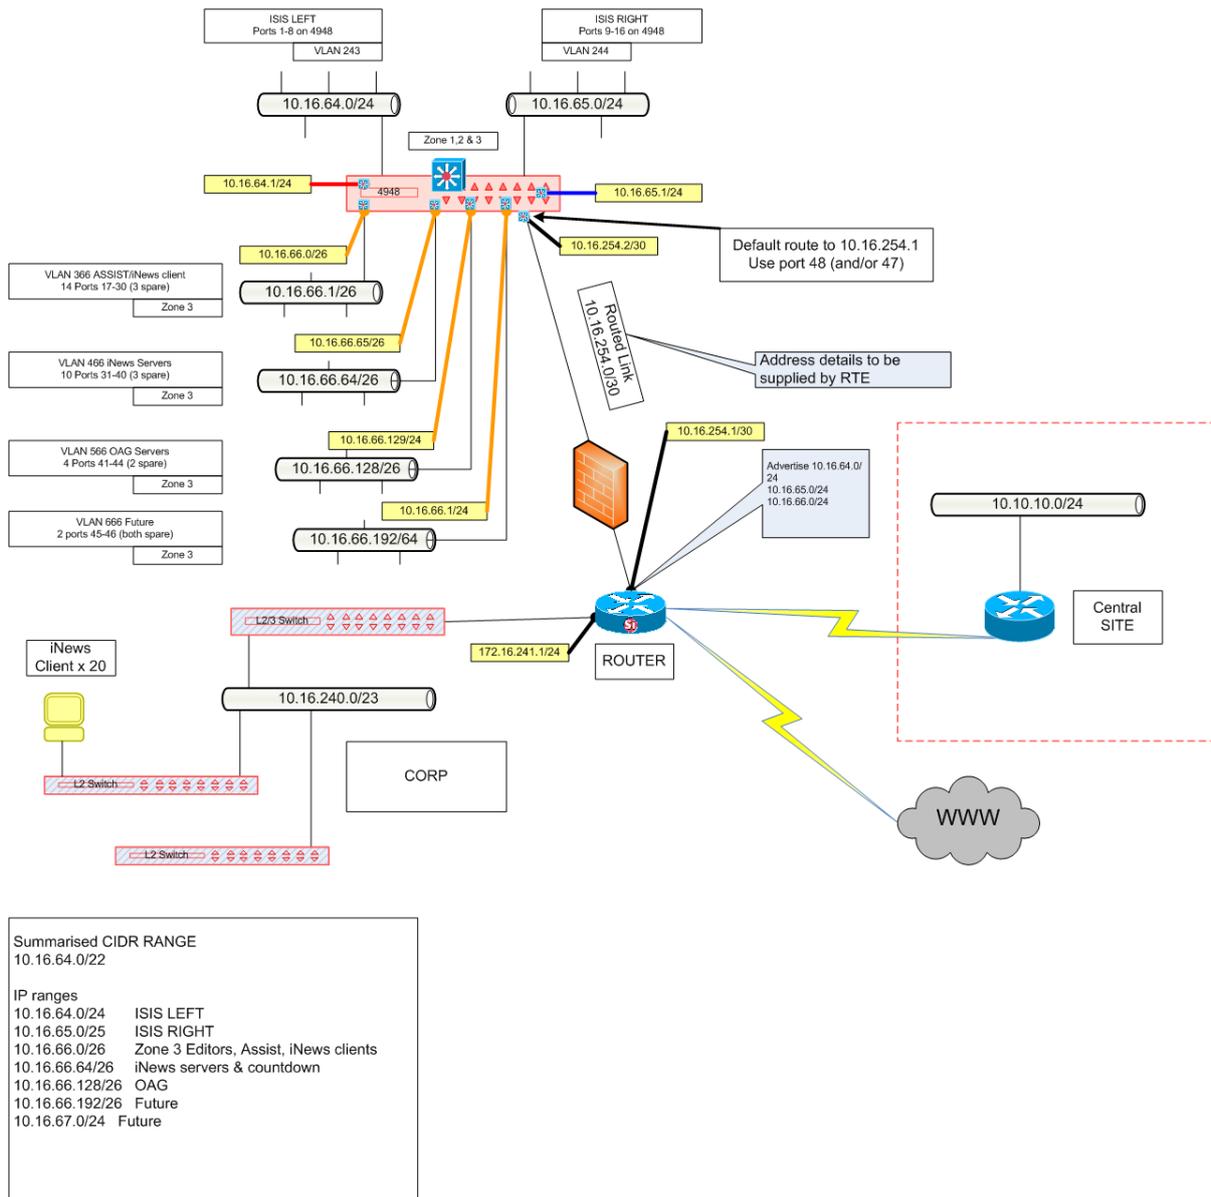
### 5.0.5 Cisco RX Core with Super X-Edge



**Figure 13 - Foundry Super X and RX MRP Core**

This network example shows a Foundry Super X VRRPe resilient border switch connection connecting into a Big Iron 40G MRP collapsed core implementation with Foundry Super X also used at the access layer.

### 5.0.6 Zone 3 Mezzanine Network conceptual diagram



**Figure 14 - High level plan of Mezzanine network structure**

This example shows a high level representation of mezzanine network where a /24 IP network as been subnetted into 4 sections to allow for granularity and the application of different firewall or access list rules.



### 5.0.7 4900M example#1 – Reference Architecture

This example shows aggregate links from ISIS with and aggregated switch interlink for resilience plus 4948 with 10G downlinks connecting clients at remote locations, such as when they are more than 90 Meters from the ISIS, but cannot deploy as Zone 4 in corporate network. This architecture has been successfully deployed at several sites.

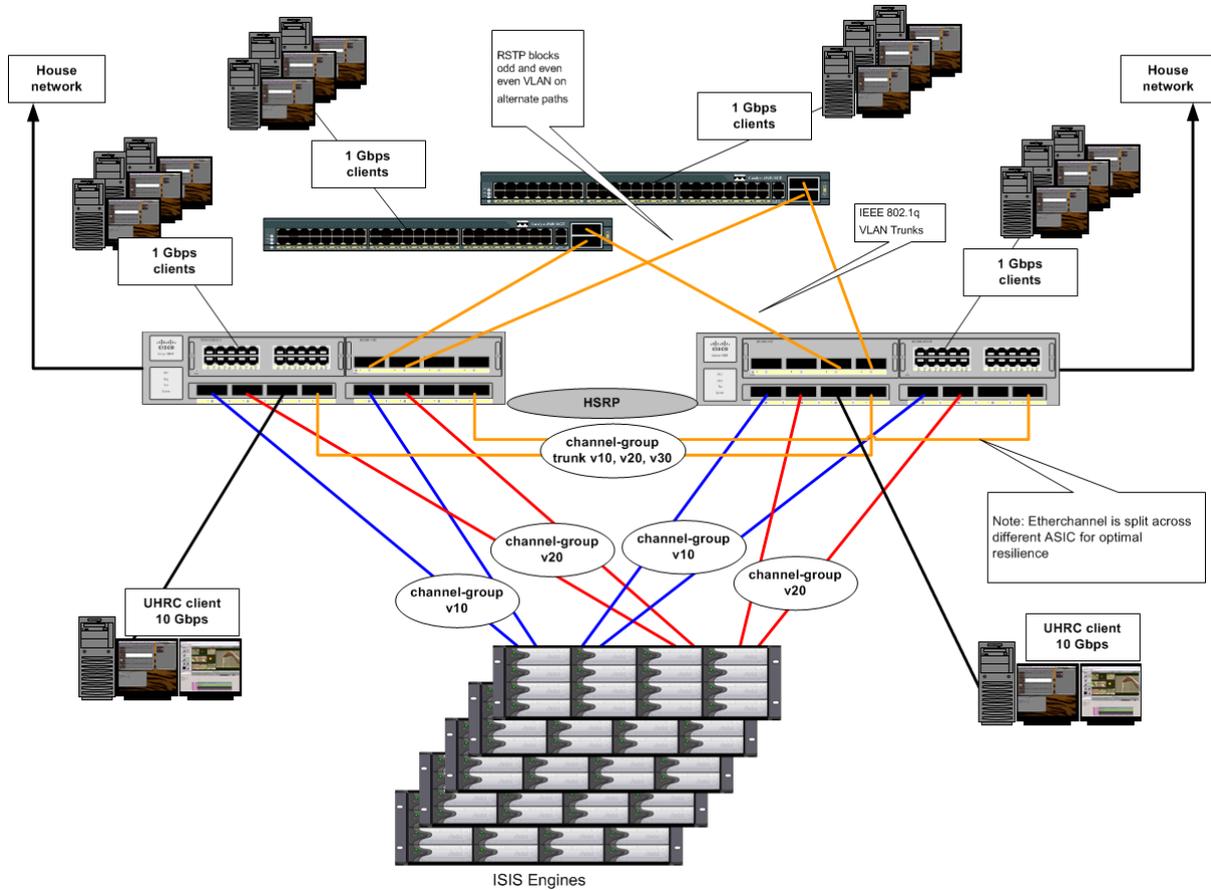


Figure 15 - Mezzanine network structure with Zone 3.1 and aggregated links

The two 4900M switches would use a resilient first hop protocol such as HSRP (or GLBP) to protect layer 3 services (see section 1.16 for details on resilient first hop protocols). The 4948 would be configured as layer 2 only.

- 

A Mezzanine network is described in more detail in Section 7.1



### 5.0.8 4900M example#2 – Reference Architecture

This example shows an extension of example#1 shown in Section 5.0.7 above, supporting 8 x C4948 cascaded switches from the C4900M pair, and using the WS-X4908 contended 2:1 10G interface module. This architecture has been successfully deployed at several sites.

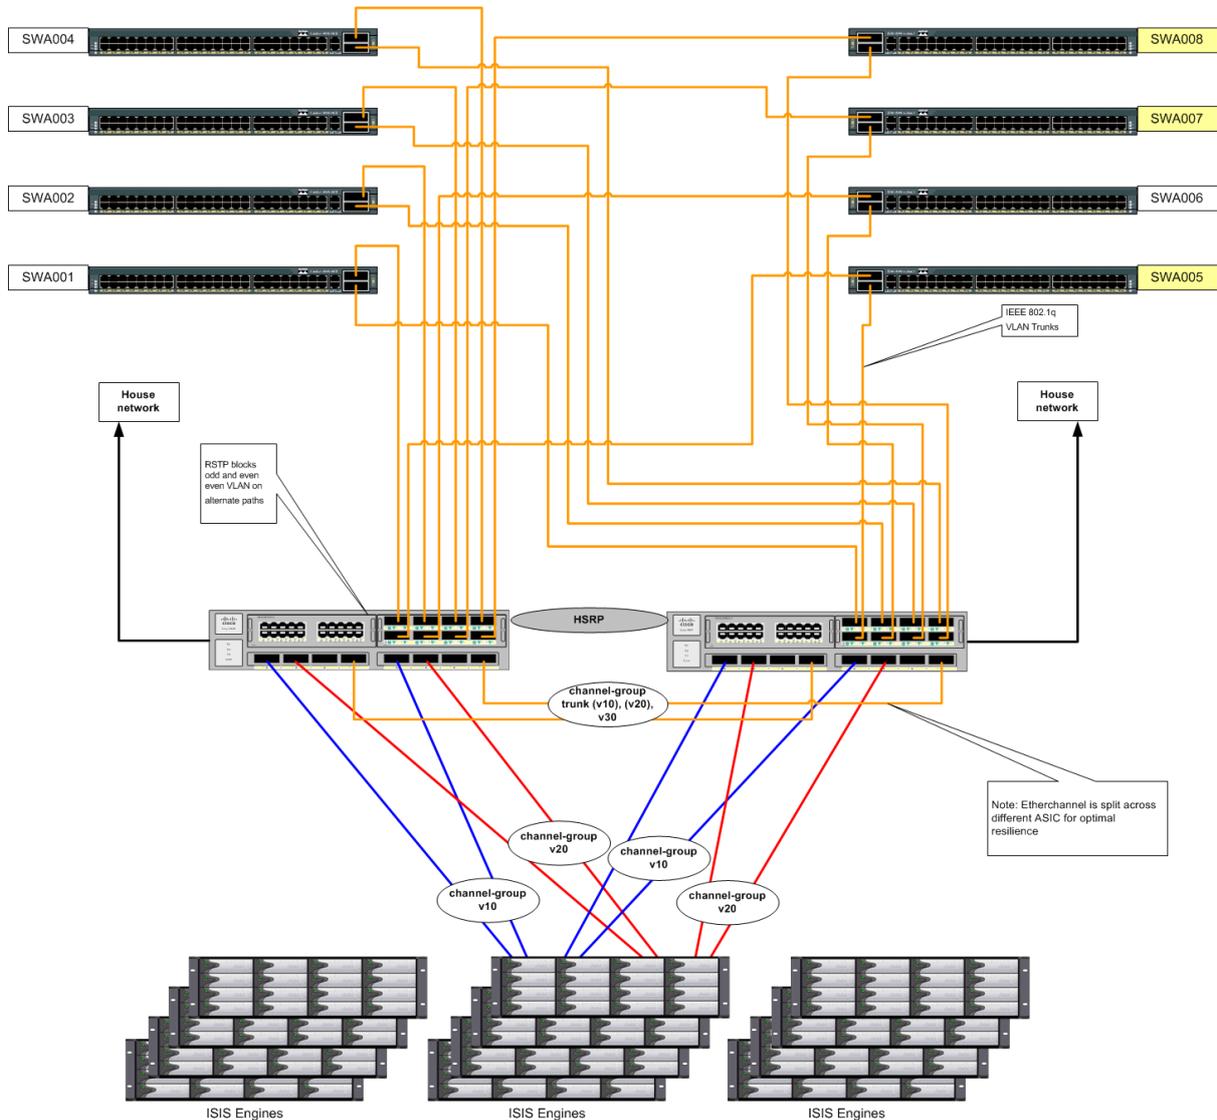


Figure 16 - Extended Mezzanine network structure with Zone 3.1 and aggregated links

Careful network design and planning will ensure that under normal operating conditions no port on 4900M ever reaches a point of contention. In most cases even in a fault situation, most links will remain uncontended as it is likely that not all client devices will be in used concurrently.

### 5.0.9 Nexus 7000 core & C4948E edge

This example shows an extension of examples 5.0.7 and 5.0.8 above using the Cisco Nexus as a core switch, Note this used a first hop resiliency protocol as VPC is not supported by ISIS 7000.



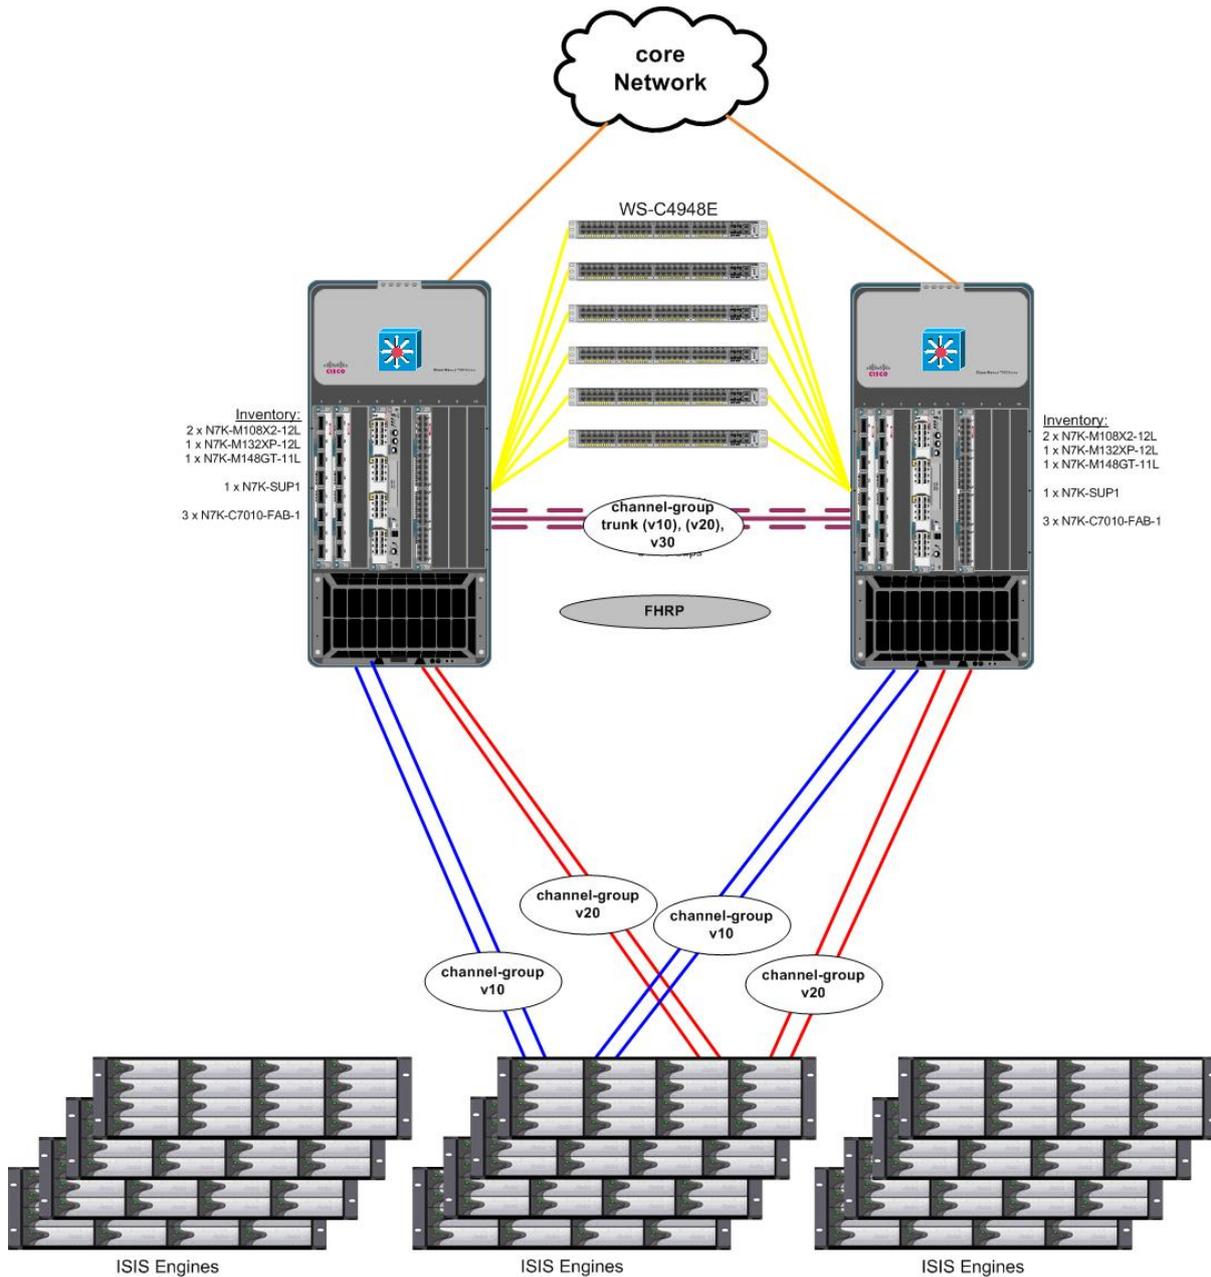


Figure 17 - Nexus 7000 core & C4948E edge

### 5.0.10 Nexus 7000 core & C4948E edge – Dual stack ISIS

When an ISIS 7000 system exceeds 12 engines, it must become a dual stack ISIS with two management domains and requires external expansion switches (EXS) as shown below. Dual Stacks can be deployed with less than 12 engines if near term expansion is likely.



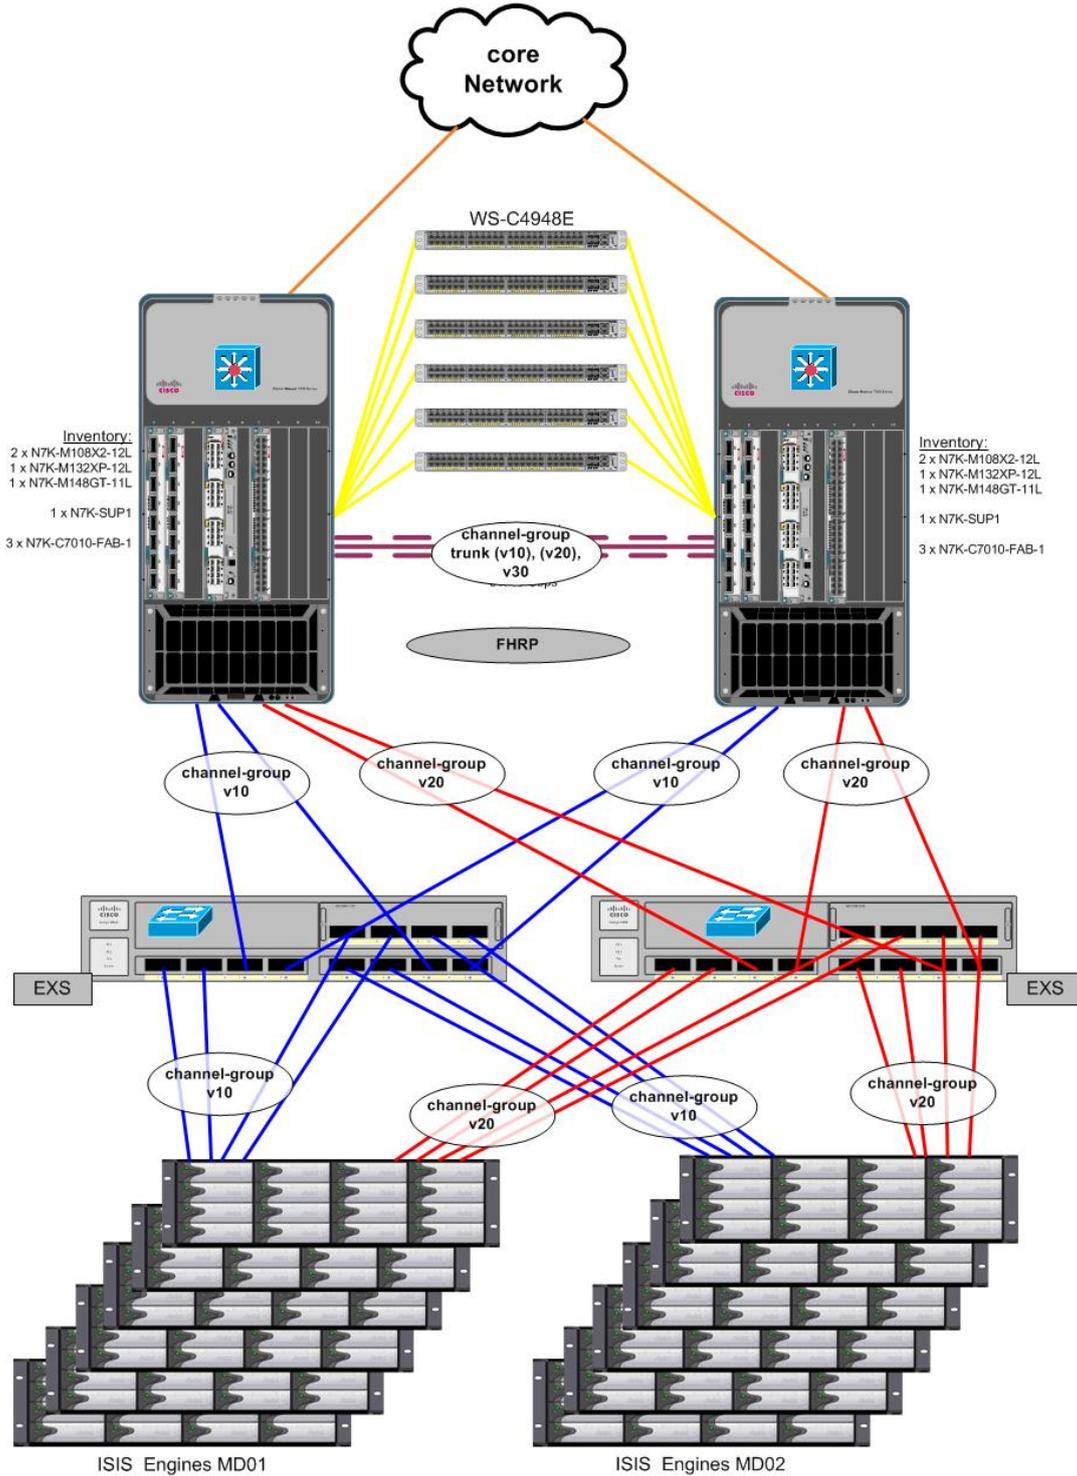


Figure 18 - Nexus 7000 core & C4948E edge Dual stack ISIS



## 5.1 Buffering

Editing clients use an oversubscription model, which mean that sufficient buffering must exist at the edge:

Hi resolution real-time video places great demands on the network infrastructure. To ensure top quality performance, by default ISIS 7000 V2.x communicates with the client using 512KB bursts of UDP traffic. This is an aggressive protocol suited to LAN environments which have high performance components. Gigabit Ethernet switches need to be able to cope with oversubscription, i.e. the ability of the 1G ports to buffer multiple 512KB bursts from the 10G interfaces. Hence this places stringent requirements on the network solution.

Network traffic from more than one interface aggregating into a single interface that does not have enough capacity, congestion is likely, and hence packet drops; this is called the aggregation problem. If joining of multiple traffic streams causes congestion on an interface, it is referred to as the confluence problem.

Avid editing clients cope with both of these issues because the oversubscription model i.e. receiving data from two ISIS storage blades concurrently. Hence the network infrastructure must also cope with this challenge.

## 5.2 Connection via IP phones – NOT RECOMMENDED

Within an IP Telephony environment it is common practice to uses an single cable to the desk and connect the PC client via the IP phone. Avid Recommend not connecting video clients via IP phones because this adds another point where the QoS can be adversely affected. Also the throughput capabilities cannot be guaranteed of this relatively low cost device (in terms of switching capabilities), plus generally the interface cards upstream of the Phone have less buffering because they are PoE based and designed for connecting lightweight clients rather than high performance real time video clients.

The worst example of this is where the IPT device has ports of different speeds and such as Gigabit Ethernet uplink to the access switch but Fast Ethernet downstream connection to the PC client. Such as device will almost certainly have insufficient buffering for the large data burst sent to an ISIS video client, and the performance will be totally unacceptable for real time video. In this configuration, the only type of Avid Interplay Production client likely to perform acceptably would be Interplay Access which only requires access to static head-frames

## 5.3 Using a dual network connection

The principle of using two different network interfaces to connect with separate networks is has been viable since ISIS 1.3. While the transitions points are not seamless, there is a “hiccup” at the time of the link failure; the editing client no longer fails when a 10G onward link fails. Also this may be done when using VMWARE, which is explained elsewhere in this document (section 7).

FROM ISIS 1.3 Readme.....Fixed Items in Avid ISIS 7000 v1.3

When using a client that is connected to Avid ISIS 7000 with dual-network ports, in a



redundant configuration, and a failure occurred with one of the network connections, there were several issues you might have experienced with the dual-network port client. This failure might have caused the Avid ISIS 7000 connection to be lost, fail, or dismount a workspace. Clients with dual-network ports also might have experienced Semaphore time-outs with 10-Gb or Hi-Gb stacking. Avid ISIS 7000 v1.3 has corrected these issues with dual-network port clients and operation now continues with a brief reaction as the client shifts from one network connection to the other.

While dual network connection to both ISIS networks is normal practice for devices such as System Director and Transfer Manager/Interplay Transfer (with a third network interface to the “outside world”), and sometimes high resolution editing platforms, using a dual network connect connection to provide connectivity to ISIS/Interplay Production system and also to a corporate network is bad practice and will lead to lots of issues and challenges. Devices which need connectivity to both resources should use a routed connection from ISIS environment to the outside network or perhaps use be implemented in Zone 3 with a mezzanine network which is explained elsewhere in this document (section 7). As a multi-homed device can present a range of issues, it may be better to use adapter teaming as explained in the section 5.4 below.

#### **5.4 Using a teamed network connection**

The principle of using two different network interfaces to connect with separate networks is A much better solution for resilient “editing type” clients (Media Composer, NewsCutter, Assist, Instinct) is to use teamed adapters, as Adapter Fault Tolerant or Switch Fault Tolerant connections. The dual port Intel Pro1000 series adapters have this functionality built in.

Implementation of the edit clients as a layer 3 device, supported by tolerant adapter teaming has been successfully deployed on several sites since 2009 and is becoming the most popular method to achieve a resilient connection which survives most fault situations.

The term “adapter teaming” has different meanings. The article below provides an excellent explanation the variations.

<http://www.intel.com/support/network/sb/cs-009747.htm>



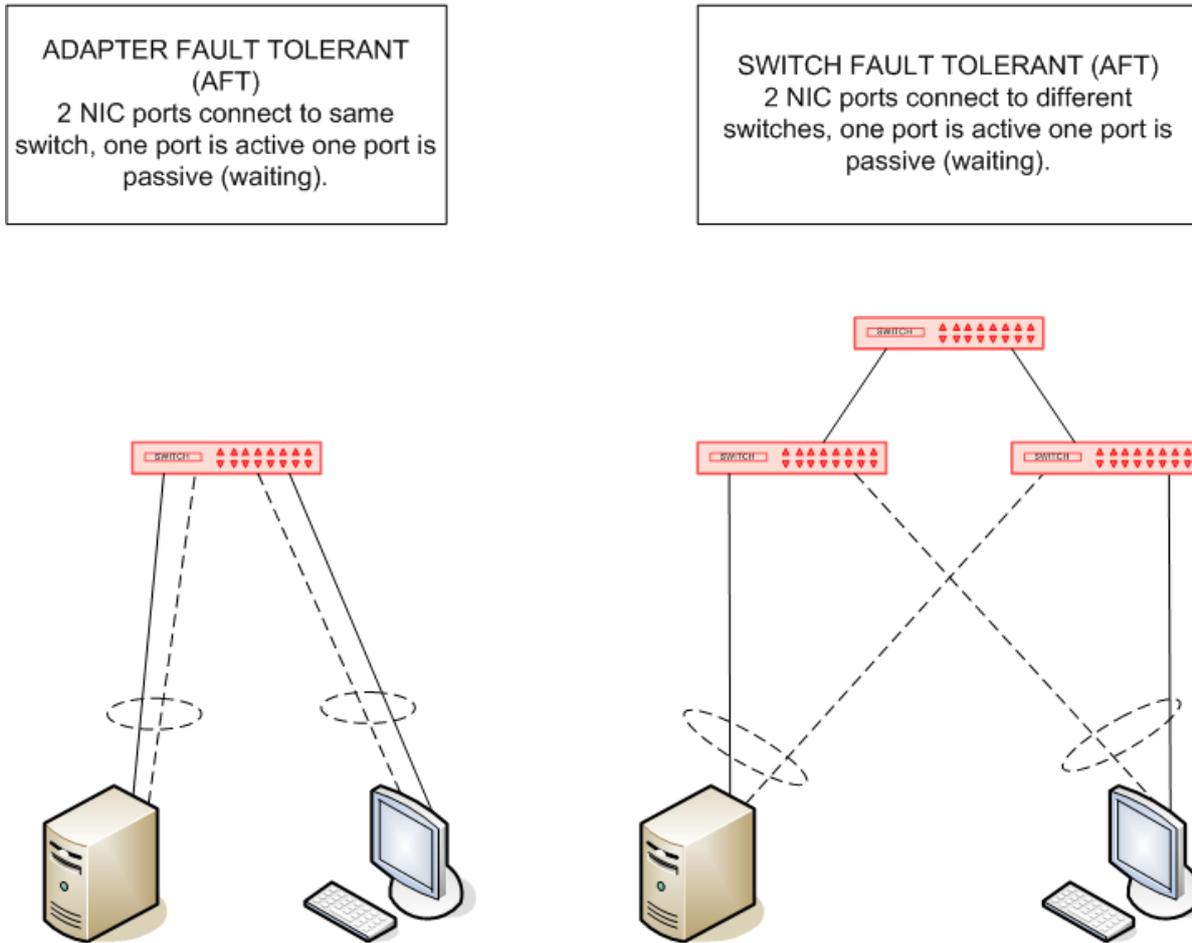


Figure 19 - AFT and SFT teaming examples

## 6.0 Firewall ISIS and Interplay Production

Many people ask the question about which firewalls are recommended or qualified for use with ISIS. Prior to Q3/2010 the short answer is NONE, this was an issue to do with ISIS and not Interplay.

NONE is a bit blunt for most customers so we have to help them understand why.

Two separate testing events of differing vendor products both in Q3/2010 with Next Generation firewall technology and 10Gigabit Ethernet interfaces that either do not re-assemble UDP datagrams by default or has very fast re-assembly have changed the previous situation (indented below).

[Initial issue V0.1 August 2008]

To the authors knowledge there are no customer successfully using a firewall for ISIS traffic. There are some who have tried and failed, and other that are trying alternative products. Also alternative approaches such as mezzanine networks of the Avid





Product Border switch, either direct or via MPLS tunnels. **\*\*NO LONGER APPLICABLE \*\***

Section 6.7 describes two successful tested next-generation firewall solutions. Sections 6.1 to 6.3 below remain explain some of the challenges posted by traditional firewall solutions.

## 6.1 First understand ISIS traffic

By default ISIS 7000 V1.x sends 256KB UDP datagrams to the client. Because the MTU size on an IP network (without Jumbo Frames) is 1500 Bytes this gets broken down by the IP stack on the ISB into 172 fragments (5 segments) which are send to the receiving client, which must re-assemble the datagram and then send up the IP stack to the application. Hence why we need a Network Interface card with lots of descriptors (1500 Byte buffers).

In ISIS 2.x the default chunk size is 512KB , which creates a total of 344 fragments/packets as 9 segments.

NOTE : Why don't we use jumbo frames? Well this is generally used for short haul TCP based server – server communications, and this is not server to server type traffic but real time video! See section 1.14.

## 6.2 Next understand Latency

ISIS client applications are latency sensitive. An editing application needs to be responsive; ISIS was designed for high speed LAN environment. When latency gets to 5ms it becomes noticeable, at 10ms it become intrusive, and at 20ms it is unpleasant to use

Some testing with NewsCutter has been done previously as part of a different products but this was based on Gigabit Ethernet MAN connection.

Latency applied	Result
0ms	System performs on test network as if locally attached
5ms	Noticeable degradation in scrubbing performance, slight delay in play function (minimal)
10ms	Particularly noticeable delay in scrubbing, 1s delay from pressing play to material playing, may not be suitable for editors
20ms	More noticeable delay in scrubbing, 2.5s delay from pressing play to material playing – this would most likely be unsuitable for editors
50ms	Unusable delay from pressing play, buffer ran out after 4-5 seconds and then started dropping frames
100ms	system will not mount ISIS workspaces, reports network errors

\*Given that the speed of light constant, 'c' is exactly 299,792,458 metres per second, the figure of 1 millisecond per 300km should be an accurate estimate for the purpose of latency calculation over distance.

Based on the tests performed with A NewsCutter editing client date 5ms is an acceptable latency; this translates to a distance of a connection of approx. 1000-1500km\* where it would be acceptable to the operator.

### 6.3 Firewall process

Normally when a firewall encounters a fragmented packet it wants to re-assemble all the fragments into a complete datagram and inspect the content from the inbound interface, when it is satisfied this is not a threat it will send the content via the outbound interface.

The first challenge for the firewall is to assemble the datagram which will be 256KB in size, this is approx 2mS on the wire, then it has to process it, then if satisfied it has to re-fragment and send on its way, so that is another 2ms, so lets call this 5ms of additional latency

The second challenge for the firewall is to re-fragment datagram in exactly the same way, which it should do under normal circumstances.

Add to this the quantity of 256KB burst per second per client, which is dependant on video resolution.

DV 25 = approx 4MB/S do that is 16 x 256KB bursts per second

DV 50 = approx 8MB/S do that is 32 x 256KB Burst per second

MPEG II Browse uncompressed audio is approx 1MB/S do that is 4 x 256KB bursts per second

Then multiply that by the number of clients, so 10 clients need 2.5MB of high speed memory available to the firewall, remember to process at high speed this need to be executed in hardware not software which would add huge amounts of latency.

This becomes an onerous task for most firewalls so the general recommendation is not to firewall ISIS video traffic. I.E. set up rules to exclude the ISIS traffic based on four criteria:

- ISIS source networks
- Client destination network
- UDP
- Ports 4000 - 4399 (version 1.0 -1.3)
- Ports 4200 - 4599 (from Version 1.4)

But firewall everything else!

But this may not work for all firewalls because even when rules for exclusion are set, the firmware in the firewall device may still complain about the number of fragments and identify this is a risk, or even continue to clog the traffic.

Another challenge is for the firewall to re-fragment in the right order, some testing proved that with a load on 1-2 clients the firewall under test could cope, but when this was increased to 10 clients the firewall under test began to change the way it applied fragmentation outgoing stream from ISIS toward the client and send packet out of order, hence corrupting the data!

Note - The port numbers for data transfer between client and ISB as well as those for incoming message traffic from the index server or ISB's are all chosen from the range 4000-4399 on a dynamic basis. That is, not all 400 ports are in use at any time, but the ports in use will vary over time over this entire range. The wide range of ports





(400 UDP ports) is used by ISIS as a tracking mechanism of the UDP segments that are exchanged between the ISIS server blade and the ISIS client.

"Note -The reason for port number changes in ISIS 1.4 refers to a conflict between Interplay Production requirements and ISIS. The Interplay Production AIF Lookup Service, the LUS (Jini/Apache River) uses port 4160 for its discovery mechanism

Prior to version 1.4 the ISIS client cycled through ports 4000 - 4400 by default even though it has only a handful of these ports open at once, it pokes holes in the firewall for all of them at boot time. Since 4160 is required for JINI then the ISIS team had to make a decision to move the ISIS ports to 4200 - 4599 to avoid any conflict."

#### **6.4 What about Interplay Production?**

Firewalling the Interplay Production traffic is not a problem, because it this is low bandwidth data packets, so do not burden the firewall, also they are not so sensitive to latency as the real time video. It is the ISIS content which is the burden for the firewall.

#### **6.5 What ports are used?**

The TCP and UDP ports used by ISIS, Interplay and WG4 are available from <http://avid.custkb.com/avid/app/selfservice/search.jsp?DocId=243397>

NOTE: THIS DOCUMENT IS NOW LEGACY, AND NO LONGER UPDATED DUE TO SEBSEQUENT DOCUMENTS RENDERING THIS A PARALLEL/DUPLICATE TASK, IT HAS BEEN DELENETD FORM THE KNOWLEDHE BASE IN EARLY 2012

:

Also see Avid Products and Network Site Preparation Guide[[373751] MAY 2011 [http://avid.force.com/pkb/articles/en\\_US/User\\_Guide/en373751](http://avid.force.com/pkb/articles/en_US/User_Guide/en373751)

Avid Products and Network Site Preparation Guide • 0130-30628-01 Rev. A • May 2011 • Created 5/16/11

#### **6.6 Why did ISIS ports used Change in ISIS V1.4?**

This was to avoid a potential port conflict with Interplay Production which used UDP port 4160 with its communications. The use of port 4160 is associated with the underlying JINI architecture.

#### **6.7 Successfully tested firewalls**

This section is a subset of the documents which can be found at on the Avid Knowledgebase at URL: <http://avid.custkb.com/avid/app/selfservice/search.jsp?DocId=244245>

In Q3/2010 there have been two successful testing event with ISIS, the newer generation of firewalls that do not automatically re-assemble fragmented UDP datagrams should not face the same challenges, as traditional firewalls experience when faces with ISIS video traffic.



### 6.7.1 Juniper SRX 3400

Testing at customer site in May and August 2010

- Juniper SRX 3400 POC testing (customer site) (Aug 2010)
  - 128MB/S successfully
  - This was the limit only because of available ISIS hardware and clients
    - Dual ISIS Engine with ISS1000, 2 clients.
  - s/w version 10.2
  - More work to do in testing
  - minimal additional latency

### 6.7.2 Cisco ASA 5500-40

Testing at Cisco Proof of Concept labs site in September 2010

- Successful PoC testing with Cisco ASA 5580-40
  - 300MB/S passed successfully
  - This was the limit only because of available ISIS hardware and clients
    - A single ISIS Engine with ISS1000, 5 clients.
  - indicative of higher bandwidth capability
  - with minimal additional latency
  - s/w version 8.3.1

### 6.7.3 Cisco FWSM for Catalyst 6500 – Limited suitability

Testing at Cisco Proof of Concept labs site in September 2010 using s/w version 4.1(2)

- with SIGNIFICANT additional latency
- limited throughput max 30MB/S
- the 24 fragment “bug” has been overcome

Prior to 4.0(2) code, the first packet of a UDP connection could not be more than 8500 bytes. With the following releases this limitation has been overcome and this product can now successfully pass AVID ISIS traffic. However this is a firewall with an old architecture and processor, therefore not powerful enough to support such an intensive task as reassembly and fragmentation of such large datagrams.

## 7.0 Security Recommendations

Avid Security Guidelines and Best Practices / Best Practice

[http://avid.force.com/pkb/articles/en\\_US/Troubleshooting/en239659](http://avid.force.com/pkb/articles/en_US/Troubleshooting/en239659)

Avid Microsoft Service Pack and Security Bulletin Addendum

[http://avid.force.com/pkb/articles/en\\_US/Troubleshooting/en239659](http://avid.force.com/pkb/articles/en_US/Troubleshooting/en239659)



## **7.1 Applying Security in Network design?**

Each customer has a different approach to security, and policies to which IT based systems should adhere to. However it must always be understood that the prime function of an ISIS/Interplay Production system is the production of video content. However, Interplay Production and iNews solutions also offer system elements which may require deployment in the corporate network or require a high level of interaction with external network services.

There are several different approaches to applying security and this section is not aimed at providing a definitive best practice, detailed configuration of network devices or authentication and authorization protocols, as such practices this will be different in each opportunity and will change with different vendor solutions. The objective is to provide some high level examples of suitable approaches.

While the ultimate solution would be to provide a separate PC for the video network and a separate PC for the corporate network, this may not be acceptable or practical due to workflow requirements or for reasons such as cost, space, power usage, or perhaps lack of available structure cabling for network connections.

### **7.1.1 Mezzanine network**

The video production network is a critical system and needs to be protected from network activities on external networks. The corporate network is also a critical system and needs protection from a system which may have different security policies and requirements such as a video production network. Real time video production traffic is not suited to firewall inspection for reasons stated elsewhere in this document (section 6).

A solution which has been deployed on many customer sites is to use a Zone 3 environment for clients which have a strong relationship with ISIS/Interplay Production i.e. need real-time video, and have a strong relationship to external networks. A good example of this is iNews Instinct client which needs access to video and needs access to an existing iNews implementation which exists in the corporate network, and also has standard iNews clients in the corporate network.

The concept of Zones in an ISIS environment are explained elsewhere in this document (section 1), and a Mezzanine network in this environment is one or more Zone 3 networks, i.e. a separate IP subnet (s), which connect via the ISIS Border switch. This allows the separation of layer 2 broadcast traffic; the use of access lists within the border switch; and the use of a perimeter firewall between the production environment and the corporate environment where there is no need for real time video. It is this “gateway” that can enforce more stringent security policies.

An example of this method is depicted below.

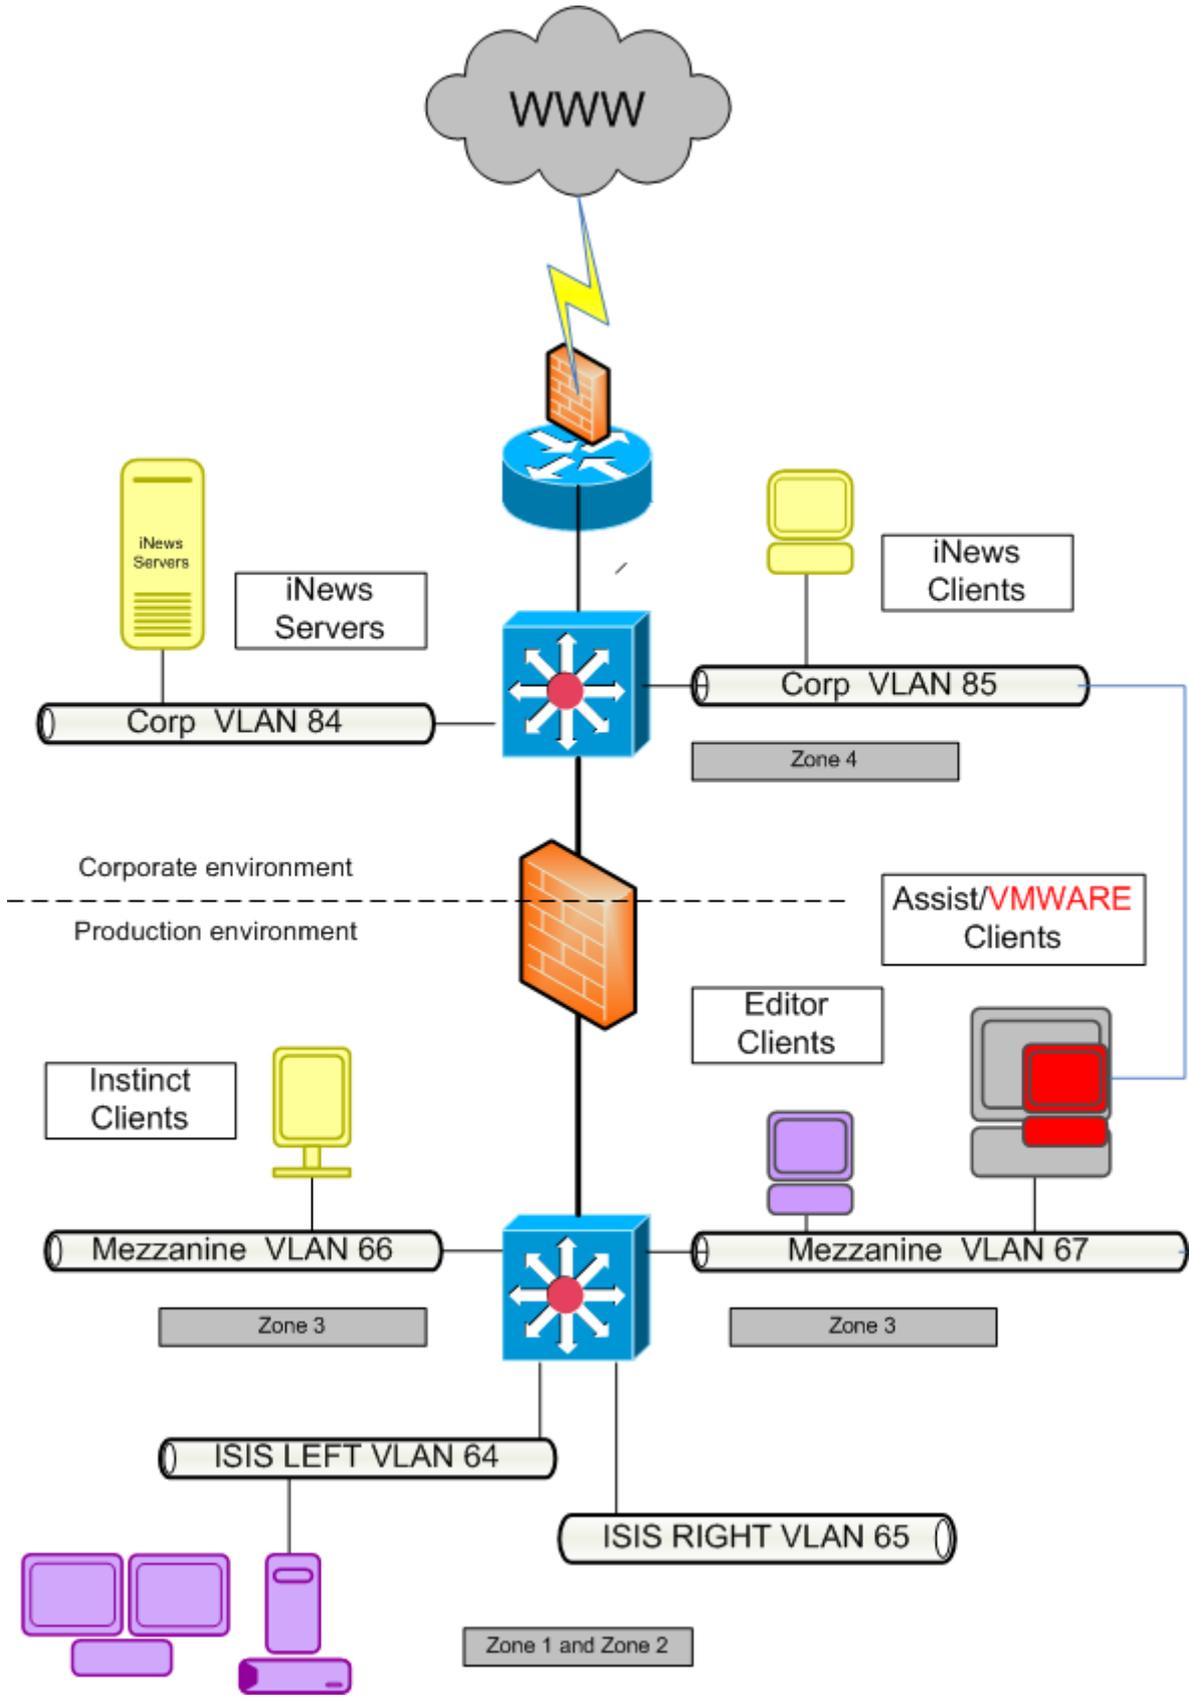


Figure 20 - Zone 3 Mezzanine Network Example



Another benefit of this design method is that providing all the necessary “heartbeat” services exist within the production environment, it can be disconnected from the corporate network, by “raising the drawbridge”, during a security situation. While there will be a loss of access to corporate services such as email, web and shared network storage, TV content can still be generated!

### 7.1.2 Using VMWARE

One way to achieve the dual PC objective with a single device is to use VMWARE to host a second virtual PC environment which is connected to an external network, via a dedicated network interface. There are several variables associated with this approach, and VMware deployment options are not addressed in detail within this section.

The Host PC would be used to run the Avid editing application and the Virtual PC would be used for access to corporate applications. The virtual PC is not active/enabled/running while the Avid editing application is in use. The host PC and the virtual PC can use different security profiles, for example the host PC may run the Avid supported Norton Antivirus solution, while the virtual PC might be running McAfee.

-----

This has been used successfully in a large scale project where the emphasis was on cost containment as well as security. Key details are given below.

The version of VMWARE USED was VMWARE ACE. This version consists of an administration console which allows the creation of create virtual machines packages. Then each package requires 1 client license when it is installed. In 2006 when the system was scoped, Admin console license was about 800€ per license and only one is required. Each client license cost around 90€

<http://www.vmware.com/products/>

Although there were significant cost savings on hardware, the virtual machine still required application licenses for the commercial applications that would exist in the virtual PC.

#### **Return on Investment?**

If the cost of a second PC suitable for the office applications was €600, and costs of the VMWARE license was € 100 approx, this reflects a saving of €500 per seat, on a 100 seat system that represents a €50K saving up front before additional desk space and power is considered an other associated services such as support and administration of additional hardware.

#### **How was it installed?**

Application side: Create a new VM computer with a 10GB virtual drive and 512MB of virtual RAM. The VM solution consisted of a full installation of windows XP in a virtual environment along with other required applications (like Office pack, PDF reader, etc...).



Then a SYSPREP was made at the end of the install, which could be applied to all target clients, via an MSI package, either via a manual install from an external drive or share, or via the network, however caution should be used on a network based install as each time it is implemented each client it will generate approx 10GB of network traffic and an MSI push of a new package could cause an extreme peak network load.

Using the SYSPREP method may not be suitable for all implementations depending on the VMware options and workflow requirements traffic which can

### **Hardware requirements?**

The computer required 1GB RAM more than the normal Avid specifications. A second network interface was required. The computer used already had a supplemental Intel Pro/1000MT network adaptor to be suitable for use with ISIS, so the on-board NIC was used by the virtual machine only and which is connected on the office VLAN, and not on the ISIS production LAN.

### **How was Editing System configured?**

Basically like every editing system. Nothing changes except for loading the VMware software.

### **Points to note:**

1. Initially this system was based on VMware ACE Version 1 but since the project concluded the customer has upgraded to VMware ACE Version 2
2. The network adaptor configuration must ensure that the adaptor for the virtual machine is not used by the physical one, and vice versa
3. Now it is possible to copy-paste text between the virtual PC and the Host PC. This feature works well in the last VMware ACE version 2.x which was not supported in VMware ACE version 1.x.
4. USB key security. If the USB key mounting is permitted in the Virtual machine it cannot be denied in the host machine.
5. Memory usage issues. On machines that did not get the 1GB RAM supplement, sometimes there are problems of memory usage since they upgrade VMware ACE in the latest version VMware ACE version 2.x.
6. The method for Upgrades, such as applying an XP service pack or application update will depend on the deployment objective. One option is to allow upgrades to be applied locally just as on a real client. Alternatively site policies might dictate that no local upgrades are permitted and a new MSI package must be created to ensure all virtual PCs run an approved image.

Depending on the setup and site policies, sharing data between host and virtual system might be allowed “internally” or might only be permitted “externally” via a mapped/ftp drive which is available to both systems and both networks.

Using an external drive/share/FTP can enhance security significantly, because real time file scanning can exist on the shared server, in fact it might be two separate servers connected with a firewall in between but which appear as the same drive letter on both the host PC and the virtual PC.

## **7.2 Internet connectivity restrictions?**

Avid has no defined policies for internet access to client devices because each customer has their own policies.

General recommendations that devices in Zone 1, 2 are denied access by use of access lists in the border switch such as the Catalyst 4900M or 4948 (and possibly elsewhere by firewall) and that devices in Zone 3, are permitted access via access lists in border switch (and possibly elsewhere by firewall).

Zone 1 and 2 devices are generally servers or craft editors so do not require internet access, and any software updating is likely to be furnished by an internal server. This status should also be applied to Zone 3 deployed Interplay Pam servers and similar/related appliances.

Zone 3 or 4 devices are likely to be Interplay PAM client application and News Editors and may require access to other corporate systems and web based applications.

Devices in Zone 4 are on the corporate LAN will normally use customer define policies.

Beyond that any use of proxies for web or other applications is entirely at the discretion of the customer policies and should not affect the performance of Avid applications.

The successful deployment of access lists are somewhat dependant on the IP addressing schema, covered elsewhere in this document. However this document does not give advice of developing access lists for use in Cisco or Brocade Layer 3 switches.

## **8.0 Network Management and Monitoring**

Network Management and Monitoring covers a wide range of discipline and mean different things to different users, to some it is about SNMP, to others it is about RMON, or perhaps link monitoring.

There are big name products such as HP OpenView, CastleRock SNMPc, Solar Winds, and independent providers such as Lortio Pro, OPManager, then there are vendor specific products such as Cisco Works and Foundry IronView which concentrate on their own networking products, and then open source solutions such as MRTG, RRDtool and Cacti products. Each provides a different and overlapping feature set.

ISIS & Interplay Production do not have a strong set of SNMP features. Being very bespoke systems, they have their own management platforms but Interplay Production can output a MIB of the system at a “moment in time”. They can send a basic trap or email alert to a Network Management System (NMS) which will invite the qualified operator to interact with their local GUI monitoring systems.



Below are details of some products that have been used on customer sites or investigated in a lab environment.

### **POWERSNMP FREE MANAGER**

A freeware, full-featured SNMP Manager application built using PowerSNMP for .NET. Discover network hosts, browse MIB trees, and analyze network requests. Suitable for lightweight to moderate management tasks.

[http://www.dart.com/psnet\\_free.aspx](http://www.dart.com/psnet_free.aspx)

OPINION: This is very good for checking out Basic SNMP functionality

### **MRTG**

The Multi Router Traffic Grapher, or just simply MRTG, is free software for monitoring and measuring the traffic load on network links. It allows the user to see traffic load on a network over time in graphical form.

It was originally developed by Tobias Oetiker and Dave Rand to monitor router traffic, but has developed into a tool that can create graphs and statistics for almost anything.

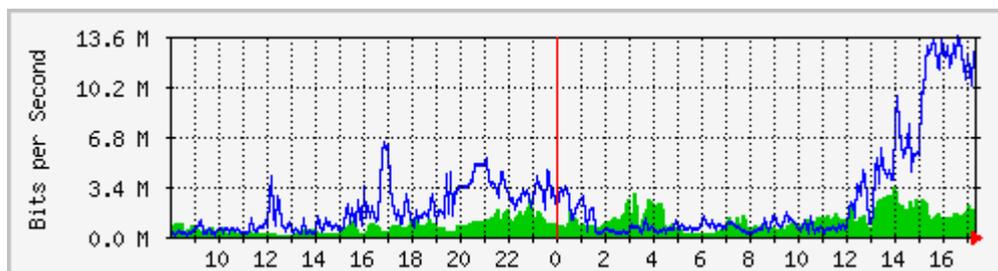
<http://en.wikipedia.org/wiki/MRTG>

<http://oss.oetiker.ch/mrtg/index.en.html>

What it does... You have a router, you want to know what it does all day long? Then MRTG is for you. It will monitor SNMP network devices and draw pretty pictures showing how much traffic has passed through each interface.

Routers are only the beginning. MRTG is being used to graph all sorts of network devices as well as everything else from weather data to vending machines.

MRTG is written in PERL and works on Unix/Linux as well as Windows and even Netware systems. MRTG is free software licensed under the Gnu GPL.



OPINION: A great tool for monitoring a small number of network interfaces but could be more challenging as the number of devices and links increase, for which Cacti is a natural progression. Fairly easy to get going on a Windows XP platform, will need Perl but documentation is good and easy to follow.

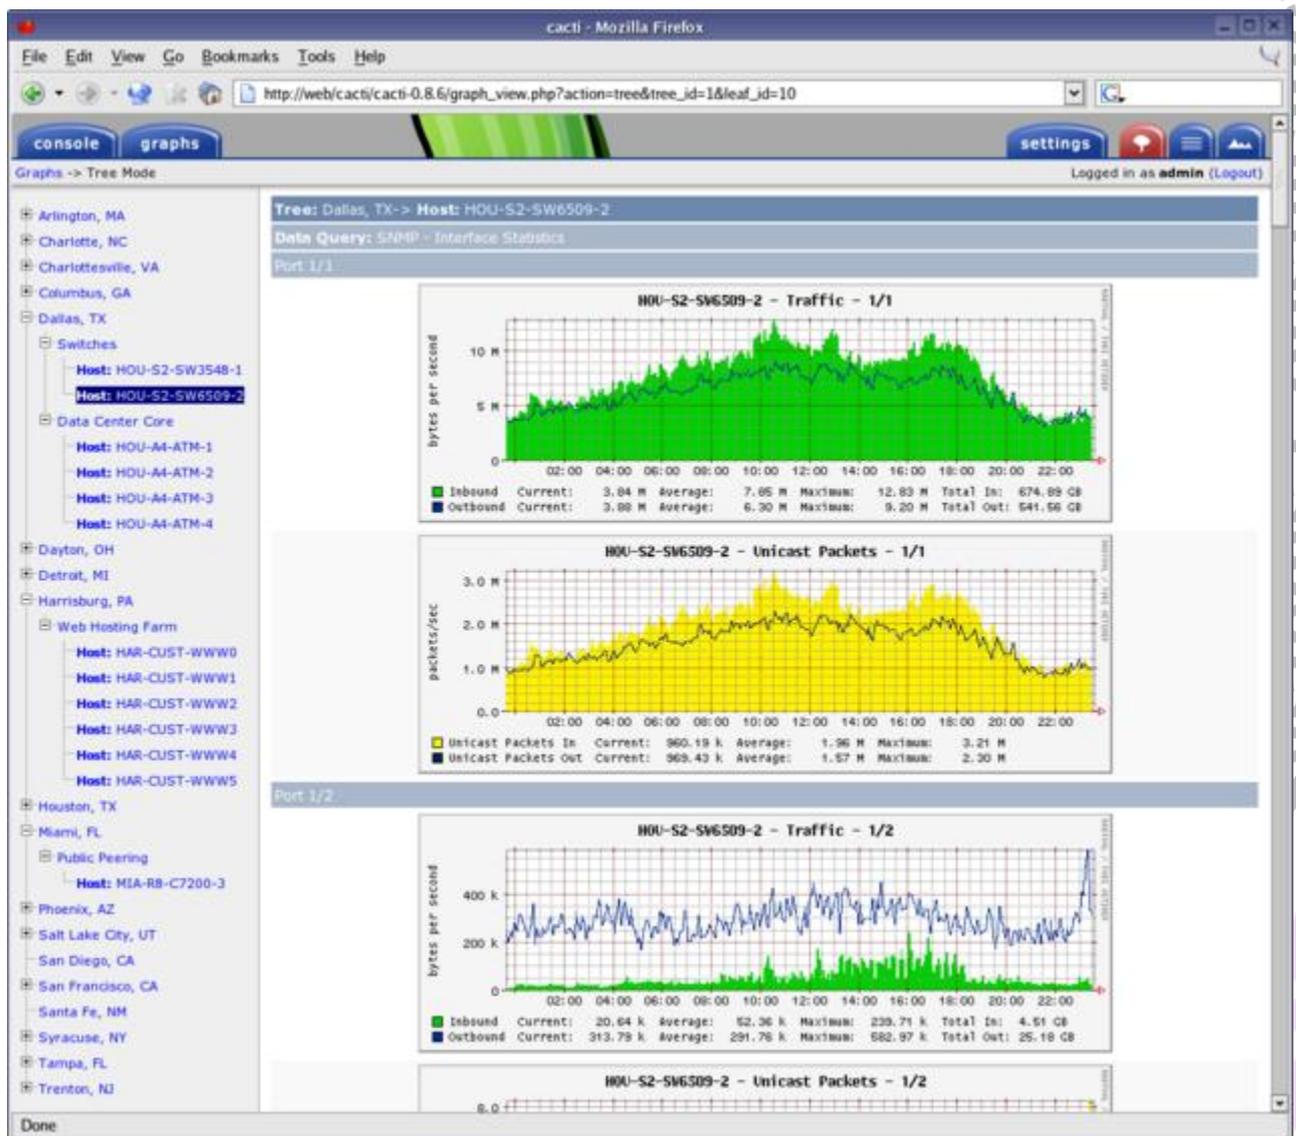
### **CACTI**

Cacti is an open source, web-based graphing tool designed as a frontend to RRDtool's data storage and graphing functionality. Cacti allows a user to poll services at predetermined intervals and graph the resulting data. It is generally used to graph time-series data like CPU load and bandwidth use. A common usage is to query network switch or router interfaces via SNMP to monitor network traffic.

[http://en.wikipedia.org/wiki/Cacti\\_\(software\)](http://en.wikipedia.org/wiki/Cacti_(software))

<http://www.cacti.net/>

Cacti is a complete network graphing solution designed to harness the power of RRDTool's data storage and graphing functionality. Cacti provides a fast poller, advanced graph templating, multiple data acquisition methods, and user management features out of the box. All of this is wrapped in an intuitive, easy to use interface that makes sense for LAN-sized installations up to complex networks with hundreds of devices.



OPINION: To some extent Cacti is a much more full featured version of MRTG which can be applied to larger systems with lots of interfaces.

**PRTG** <http://www.paessler.com/prtg>

PRTG Traffic Grapher is the perfect small tool for quick and easy bandwidth monitoring via SNMP, NetFlow and packet sniffing. It installs in a few minutes and provides current and historic bandwidth usage data for your network

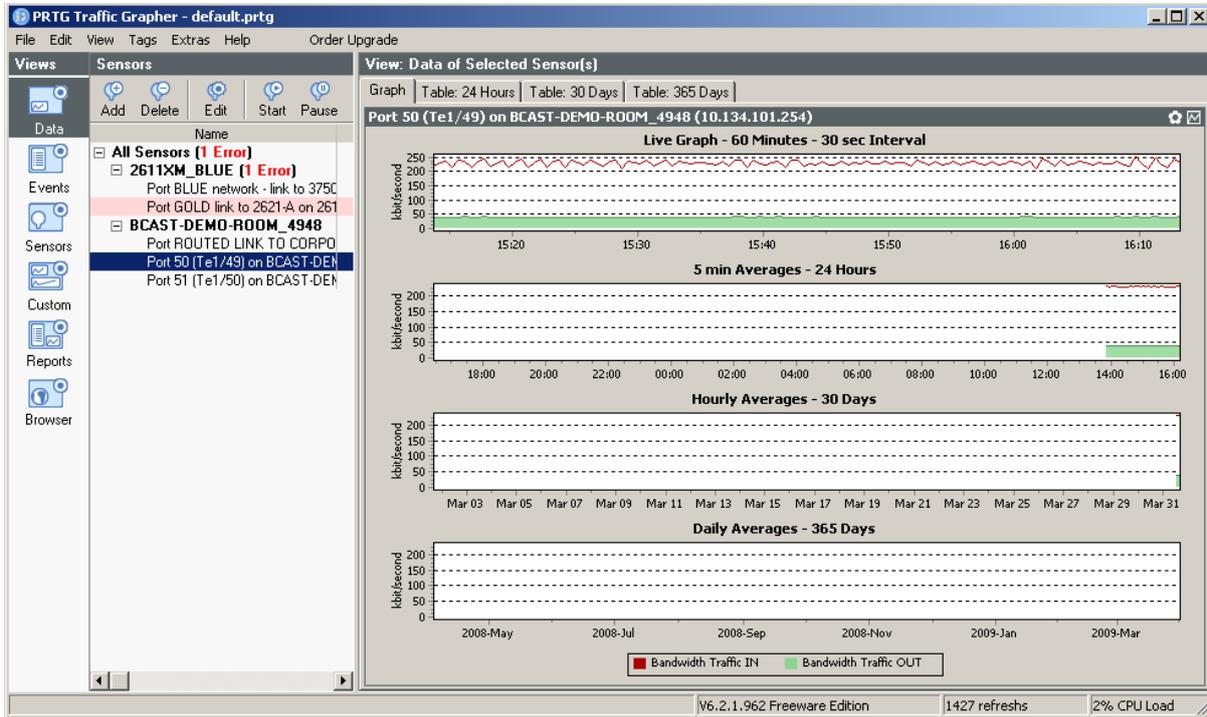
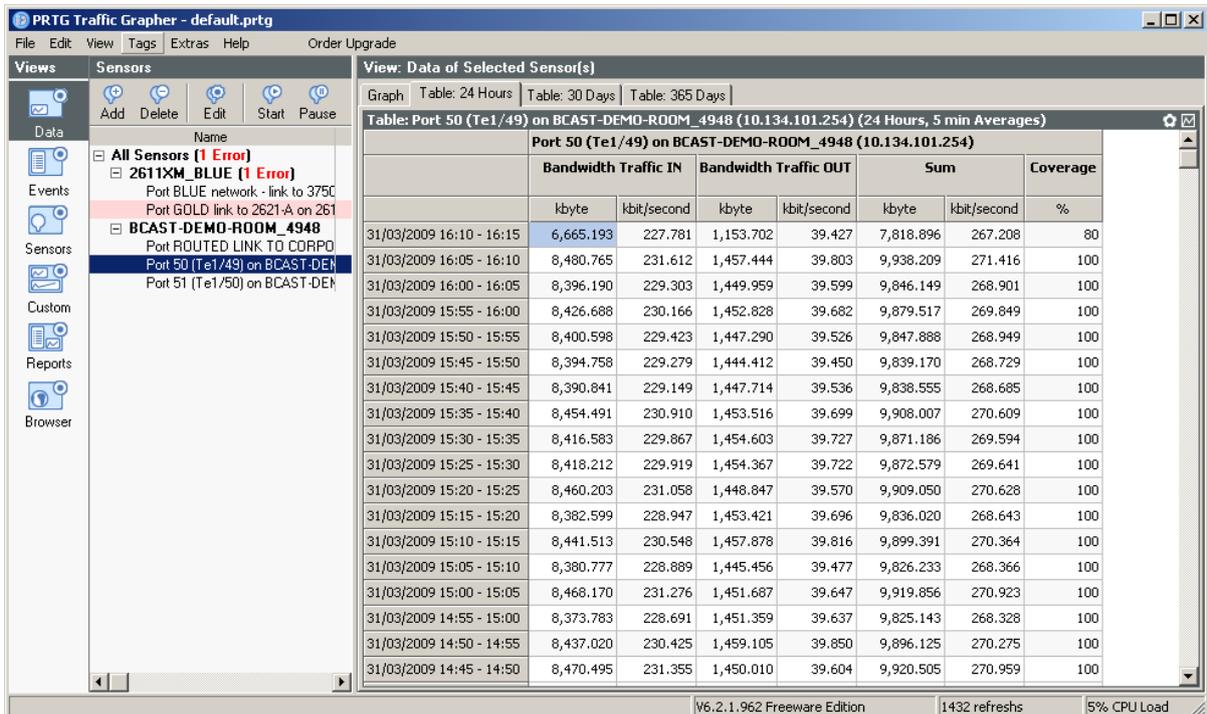



Table: Port 50 (Te1/49) on BCAST-DEMO-ROOM_4948 (10.134.101.254) (24 Hours, 5 min Averages)							
	Bandwidth Traffic IN		Bandwidth Traffic OUT		Sum		Coverage
	kbyte	kbit/second	kbyte	kbit/second	kbyte	kbit/second	%
31/03/2009 16:10 - 16:15	6,665.193	227.781	1,153.702	39.427	7,818.896	267.208	80
31/03/2009 16:05 - 16:10	8,460.765	231.612	1,457.444	39.803	9,938.209	271.416	100
31/03/2009 16:00 - 16:05	8,396.190	229.303	1,449.959	39.599	9,846.149	268.901	100
31/03/2009 15:55 - 16:00	8,426.688	230.166	1,452.828	39.682	9,879.517	269.849	100
31/03/2009 15:50 - 15:55	8,400.598	229.423	1,447.290	39.526	9,847.888	268.949	100
31/03/2009 15:45 - 15:50	8,394.758	229.279	1,444.412	39.450	9,839.170	268.729	100
31/03/2009 15:40 - 15:45	8,390.841	229.149	1,447.714	39.536	9,838.555	268.685	100
31/03/2009 15:35 - 15:40	8,454.491	230.910	1,453.516	39.699	9,908.007	270.609	100
31/03/2009 15:30 - 15:35	8,416.583	229.867	1,454.603	39.727	9,871.186	269.594	100
31/03/2009 15:25 - 15:30	8,418.212	229.919	1,454.367	39.722	9,872.579	269.641	100
31/03/2009 15:20 - 15:25	8,460.203	231.058	1,448.847	39.570	9,909.050	270.628	100
31/03/2009 15:15 - 15:20	8,382.599	228.947	1,453.421	39.696	9,836.020	268.643	100
31/03/2009 15:10 - 15:15	8,441.513	230.548	1,457.878	39.816	9,899.391	270.364	100
31/03/2009 15:05 - 15:10	8,380.777	228.889	1,445.456	39.477	9,826.233	268.366	100
31/03/2009 15:00 - 15:05	8,468.170	231.276	1,451.687	39.647	9,919.856	270.923	100
31/03/2009 14:55 - 15:00	8,373.783	228.691	1,451.359	39.637	9,825.143	268.328	100
31/03/2009 14:50 - 14:55	8,437.020	230.425	1,459.105	39.850	9,896.125	270.275	100
31/03/2009 14:45 - 14:50	8,470.495	231.355	1,450.010	39.604	9,920.505	270.959	100



PRTG Network Monitor is the powerful network monitoring solution from Paessler AG. It monitors your network using a whole range of technologies and assures the availability of network components and measures traffic and usage. It saves costs by avoiding outages, optimizing connections, saving time and controlling service level agreements (SLAs).

A freeware version of both products are available which monitor up to 10 interfaces which will be fine for small installations, to look at 10G link utilization, and selected 1Gigabit Ethernet interfaces.

### **CASTLEROCK**

Castle Rock Computing SNMPC, is a secure distributed network management system which delivers proactive real-time monitoring for your entire network infrastructure. Advanced product features and legendary ease of use have led to over 120,000 network managers trusting SNMPC to monitor their mission critical networks.

#### Key Product Features:

- Monitors SNMP devices, WAN Links, Servers and Applications
- Supports SNMP v1, v2c and secure SNMP v3
- Scalable, Distributed Architecture
- Email/Pager Event Notification
- Vendor Independent - Manages any SNMP device from any vendor
- Key Network Metrics (e.g. Utilization)
- Automatic WEB & Printed Trend Reports
- Live/Standby Servers with automatic failover
- Automatic Baseline Alarms
- Runs as Windows Service
- Remote Console & JAVA Access
- Real-time MIB Displays
- Automated Network Discovery
- Programming & Scripting Interfaces

<http://www.castlerock.com/products/snmpc/default.php>

### **LORIOT PRO**

[http://www.loriotpro.com/Products/Features/Summary\\_Free\\_Edition\\_EN.php](http://www.loriotpro.com/Products/Features/Summary_Free_Edition_EN.php)

The freeware edition of the LorientPro software gives you access to a powerful graphical SNMP manager. This free Windows-based software helps you to access SNMP devices, to create IP network maps and directory maps, manage IP routers, to analyze SNMP requests, to perform SNMP get and set requests, to compile MIB files, to browse MIB tree, to receive events and SNMP trap, to discover networks and hosts. (FREEWARE EDITION LIMITED TO 10 HOSTS)

OPINION: Seems very comprehensive, but then there is so much, it will need good amount of time investing to get the best out of it.

### **OPMANAGER**



OpManager is a complete, end-to-end Network & IT infrastructure monitoring platform that offers advanced fault and performance management across WAN, VoIP services, network devices, servers, applications, databases and other IT infrastructure such as printers, UPS etc. (FREEWARE EDITION LIMITED TO 10 HOSTS)

<http://manageengine.adventnet.com/products/opmanager/>

OPINION: This is very good package for trap management and basic device management but getting the port throughput traces proved challenging. Only tried on network switches. Probably need more time investing to get the best out of it

## 9.0 DNxHD in Zone 3 and 4

The deployment of DNxHD edit clients in Zone 4 was not a workflow significantly deployed in ISIS 1.x solutions. The ISIS 2.0 platform (initially release DEC 2008 ) makes this a scalable solution.

A Zone 4 deployment relies on the customer provided network delivering suitable bandwidth with a latency figure of less than 5mS (measured with PATHDIAG).

Avid have validated this solution in the testing described below, using a multiple switch hops.

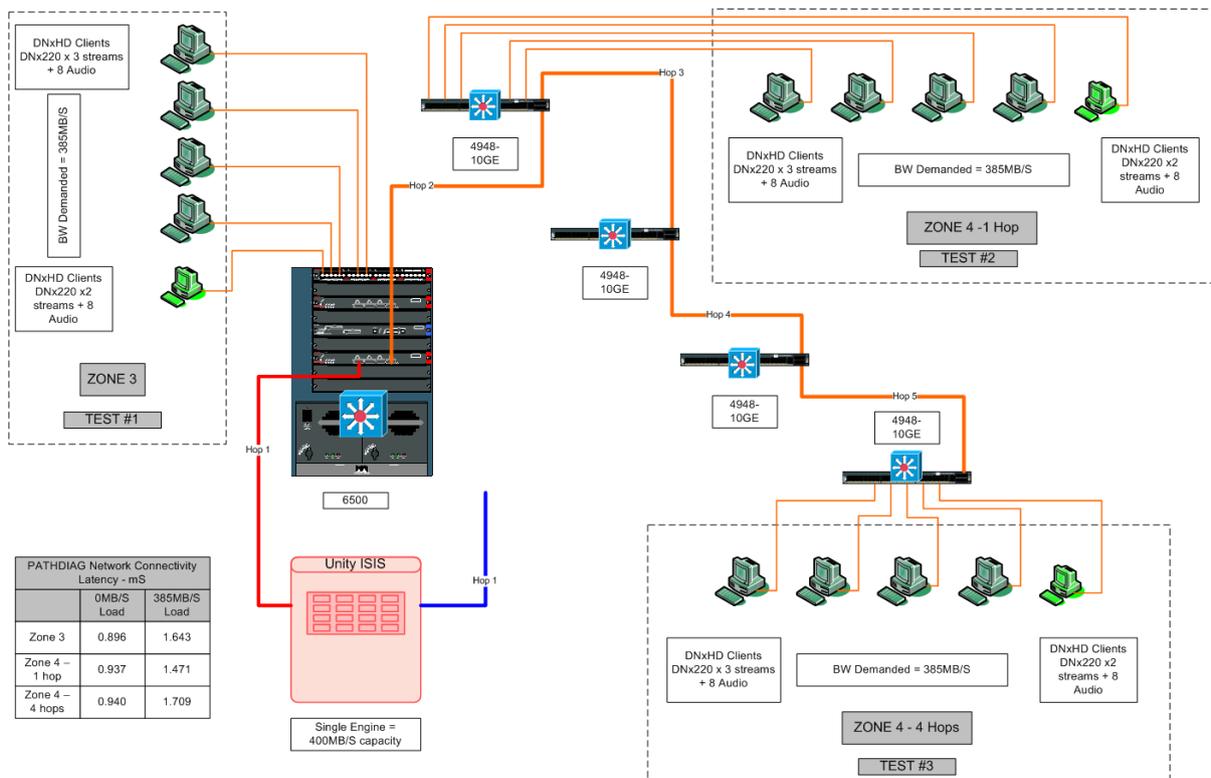


Figure 21 - DNxHD in Zone 3 and 4

## 9.1 Test Setup:

A single ISIS 2.0 Engine was connected via a single 10G path to a Cisco Catalyst 6500 switch and then via four 4948-10GE switches. A single engine is rated to provide 400MB/S.

The standard “2 second cut” based sequence was used as test media which places a significant burden on the ISIS and the Editor. Gigabit Ethernet connected editing clients were configured to pull a minimum of stream 2 x DNX220 streams with 8 Audio streams. Total load of Zone 4 clients was 385MB/S, which exceeds projected Bandwidth demand of the CUSTOMER work flow. (16 x DNx120 streams = 248MB/S)

### 9.1.1 Test Equipment

1 SR2500 System Director, Windows XP embedded 32-bit & service pack 2 with ISIS 2.0.1 gold build

1 ISIS chassis with 2 ISS2 switches & 16 2-TB ISBs with ISIS 2.0.1 gold build; single 10 Gb Ethernet connection (no link aggregation), mirrored workspace & 512KB chunk size

1 Cisco 6509 & 1 WS-X6748-GE-TX module for 1 Gb client connections & 2 WS-X6704-10GE modules for 10 Gb Ethernet connections between switches (no link aggregation), with IOS 12.2(17d)SXB7 & EIGRP enabled for routing; zone 3 switch

4 x Cisco 4948's with IOS 12.2(25).EWA10 & EIGRP enabled for routing; zone 4 switches

5 HP XW8600's, Vista 64-bit clients & service pack 1, Intel Pro 1000 PT dual port network adapters with driver 9.12.30.0 (part of 13.3 package) with ISIS 2.0.1 gold build & Media Composer 3.5 gold build, (no Interplay)

### 9.1.2 Test Results

Test 1

Playback 4 clients @ DNxHD220 3 video & 8 audio tracks & 1 client @ DNxHD220 2 video & 8 audio tracks on Zone 3 switch – Pass; 0 skipped frames, fill rate = 30 frames/sec for 3 video track clients & 50 frames/sec for 2 video track client

Test 2

Playback 4 clients @ DNxHD220 3 video & 8 audio tracks & 1 client @ DNxHD220 2 video & 8 audio tracks on 1st Zone 4 switch (1 hop away from ISIS) – Pass; 0 skipped frames, fill rate = 30 frames/sec for 3 video track clients & 50 frames/sec for 2 video track client

Test 3

Playback 4 clients @ DNxHD220 3 video & 8 audio tracks & 1 client @ DNxHD220 2 video & 8 audio tracks on 4th Zone 4 switch (4 hops away ISIS) – Pass; 0 skipped frames, fill rate = 30 frames/sec for 3 video track clients & 50 frames/sec for 2 video track client

## 9.2 Test Summary

PATHDIAG Network Connectivity Latency – milliseconds								
	0 MB/S Load				385 MB/ Load			
	LO	HI	AVG	DELTA Z3	LO	HI	AVG	DELTA Z3
Zone 3	0.806	1.064	0.896		0.631	5.445	1.643	
Zone 4, 1 hop	0.851	1.080	0.937	0.041	0.611	5.470	1.471	(0.172)
Zone 4, 4 hops	0.847	2.866	0.940	0.044	0.656	5.45	1.709	0.066

The table above shows the latency measurements taken using the PATHDIAG network connectivity function which sends and 8192 PING targeted, in this testing, to the ISB blade in the ISIS engine from an Avid client in each of the test scenarios. The average difference based on differences between Zone 3 and Zone 4-4 hops is Less that 0.1 millisecond with or without loading, the majority of the increase in PING time being due to loading in the ISIS.



The screenshot shows the Avid PathDiag interface. The 'Settings' section includes 'Test: Custom Test', 'Duration (Minutes): 2', and 'Path to Test: z\''. The 'Control' section shows 'Elapsed Time: 02:03' and 'Time Remaining: -'. A message states 'Test stopped Wed March 11, 2009 11:02:27'. The 'Results' section displays 'Network: Connectivity Test' with 'IOs: 121' and 'Errors: 0'. Below this is a table of test results:

Source	Target	Sent	Recv	Lo RTT (ms)	Hi RTT (ms)	Avg RTT (ms)	Lo Hops	Hi Hops	Avg Hops	% Lost
168.30.214	192.168.10.25 (192.168.10.25)	121	121	0.806037	1.064730	0.896368	1	1	1	0

The averages were taken based on 120 pings, over 2 minutes and shown in the screen shots above.

This proves that the desired goal should be easily achievable on a suitable enabled customer network. However network consultancy services from Avid should always be engaged when this is a project deliverable



## APPENDIX A. How to integrate Interplay Production Engine in trusted domain environments

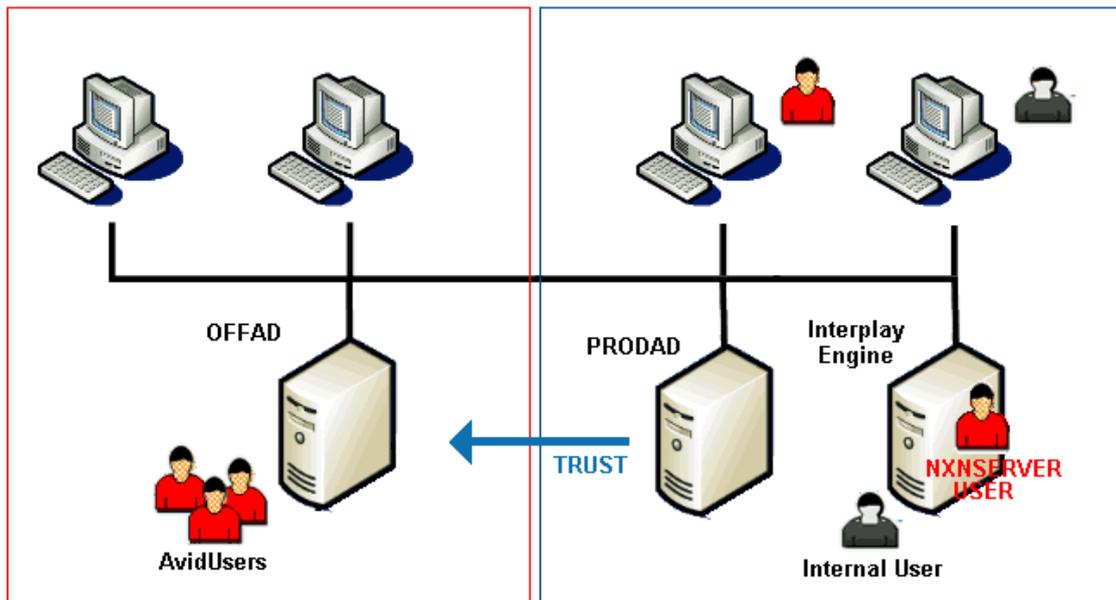
Thanks to Ralf for this content

### Scope

This document explains how to integrate the Interplay Engine server into a trusted domain consisting on an office and production domain. So users within the office domain can be used for authentication on the Interplay Production systems.

### Scenario

1. Existing Office domain (naming OFFAD)
2. Existing Production domain (naming PRODAD)
3. Outgoing trust from PRODAD to OFFAD
4. User group (naming AvidUsers) within the OFFAD which should be used for authentication on Interplay Production (using the import feature within the admin client)



### Configuration requirements

1. Administrator username and password from OFFAD
2. All machines joined their respective domains, e.g. machines within the PRODAD Network, the PRODAD domain etc. The Interplay Engine machine(s) should also join the PRODAD domain. In case of a cluster configuration the cluster service account should be created within the PRODAD domain.
3. Create NXNServer execution user account within the OFFAD domain





**Interplay Production users within the PRODAD are not able to authenticate on the Interplay Engine after the changes, they must be created internal or imported by other authentication providers (Unity, LDAP)**

## **Configuration steps**

### **A. Standalone Interplay Production Engine**

1. Specify the NXNServer execution user account during installation of the Interplay Engine or use the change tool. The installer and the tool will add the user automatically to the local administrator group.
2. verify proper change by opening the taskmanager process list and check the nxnserver process owner.
3. Import the AvidUsers group within the Interplay Production Administrator client

### **B. Clustered Interplay Production Engine**

1. Create a cluster service user account within the OFFAD. Add the user to the following security policy settings:

- Act as part of the Operating System
- Adjust memory quotas
- Backup operation
- Increase scheduling priority
- Log on as a service
- Restore files
- Manage audit
- Impersonate a user
- Debug programs

2. Specify the NXNServer execution user account during installation of the Interplay Engine or use the change tool. The installer and the tool will add the user automatically to the local administrator group. Verify the registry contains the Execution user settings correctly!
3. Reboot cluster and change the cluster service account within the service control panel of the cluster service user account within OFFAD.
4. Reboot and verify proper startup of all cluster resources
5. verify proper change by opening the taskmanager process list and check the nxnserver process owner.
6. Import the AvidUsers group within the Interplay Production Administrator client



## APPENDIX B. Switch configuration tips & good practices

This is not an exhaustive list of best practices, as these may vary from site to site based on customer policies and procedures and security requirements.

### ***B.1 Document your configs with DESCRIPTIONS***

Use the description command to apply simple description to interfaces and VLANs such as AIRSPEED, MEDIA INDEXER, INTERPLAY ENGINE

In Foundry the name or port-name command is used.

#### **A simple description**

```
conf t
interface g1/1
description AIRSPEED
```

#### **CISCO**

```
PWDEMO-Cisco4948#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PWDEMO-Cisco4948 (config)#int t1/49
PWDEMO-Cisco4948 (config-if)#desc CONNECTION TO ISIS VLAN LEFT
PWDEMO-Cisco4948 (config-if)#int t1/50
PWDEMO-Cisco4948 (config-if)#desc CONNECTION TO ISIS VLAN RIGHT
PWDEMO-Cisco4948 (config-if)#exit
PWDEMO-Cisco4948 (config)#exit
PWDEMO-Cisco4948#
```

#### **FOUNDRY**

```
telnet@ETSG_SWITCH#conf t
telnet@ETSG_SWITCH(config)#
telnet@ETSG_SWITCH(config)#int e 22
telnet@ETSG_SWITCH(config-if-e1000-22)#port-name SYSTEM UNDER TEST D-
SHEPHARD
telnet@ETSG_SWITCH(config-if-e1000-22)#exit
telnet@ETSG_SWITCH(config)#
```

### ***B.2 Setting Spanning tree to Rapid Spanning Tree***

#### **CISCO**

```
PWDEMO-Cisco4948#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PWDEMO-Cisco4948 (config)#span
PWDEMO-Cisco4948 (config)#spanning-tree mode rap
PWDEMO-Cisco4948 (config)#spanning-tree mode rapid-pvst
PWDEMO-Cisco4948 (config)#exit
PWDEMO-Cisco4948#
PWDEMO-Cisco4948#
```

#### **FOUNDRY**

```
BigIron(config)# spanning-tree 802-1w
```



or

```
BigIron(config)# vlan 10  
BigIron(config-vlan-10)# spanning-tree rstp
```

## B2.1 Spanning tree cost

When costing up links this can be done on a per link basis or on a per VLAN basis across a given link

### PER LINK/INTERFACE

```
PWDEMO-Cisco4948(config-if)#spanning-tree cost 10
```

### PER VLAN, PER LINK.INTERFACE

```
PWDEMO-Cisco4948(config-if)#spanning-tree vlan 10,20 cost 10
```

The value of 10 is appropriate when the short method is used spanning-tree path cost. When using the long method for spanning-tree path cost, a value of 5,000 is appropriate. This value is chosen so as not to reflect a pre-define value.

## B.2.2 Spanning Cost type

Spanning tree cost can be Long or Short, different value applied as per the links below

<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.1/12.1e/command/reference/S1.html#wp1029022>

### Usage Guidelines

This command applies to all the spanning tree instances on the switch.

The **long** path cost calculation method uses all the 32 bits for path cost calculation and yields values in the range of 1 through 200,000,000.

The **short** path cost calculation method (16 bits) yields values in the range of 1 through 65,535.

### Examples

This example shows how to set the path cost calculation method to long:

```
Switch(config)# spanning-tree pathcost method long  
Switch(config)#
```



This example shows how to set the path cost calculation method to short:

```
Switch(config)# spanning-tree pathcost method short
Switch(config)
```

[http://en.wikipedia.org/wiki/Spanning\\_Tree\\_Protocol#Data\\_rate\\_and\\_STP\\_path\\_cost](http://en.wikipedia.org/wiki/Spanning_Tree_Protocol#Data_rate_and_STP_path_cost)

## Data rate and STP path cost

The table below shows the default cost of an interface for a given data rate.

Data rate	STP Cost (802.1D-1998)	STP Cost (802.1t-2001)
4 Mbit/s	250	5,000,000
10 Mbit/s	100	2,000,000
16 Mbit/s	62	1,250,000
100 Mbit/s	19	200,000
1 Gbit/s	4	20,000
2 Gbit/s	3	10,000
10 Gbit/s	2	2,000

### B.3 SET primary switch as STP master root primary

Spanning tree settings must be done for each VLAN

```
PWEMO-Cisco4948#conf t
PWDEMO-Cisco4948(config)#spanning-tree vlan [NUM] root primary
PWDEMO-Cisco4948(config)#exit
PWDEMO-Cisco4948#
```

The command that actually gets entered in to Cisco running config will look like this `spanning-tree priority 24576` but the number will be change depending on the VLAN numbers.

#### FOUNDRY

```
BigIron(config)#spanning-tree priority 24576
```

Syntax: [no] spanning-tree [forward-delay <value>] | [hello-time <value>] | [maximum-age <value>] | [priority <value>]

Here is the syntax for port STP parameters.

Syntax: [no] spanning-tree ethernet | pos <portnum> path-cost <value> | priority <value>

priority: Possible values: 1 – 65535. Default is 32768. A higher numerical value means a lower priority;

thus, the highest priority is 0.



## **B.4. SET secondary switch as STP root secondary**

Spanning tree settings must be done for each VLAN

```
PWEMO-Cisco4948#conf t
PWDEMO-Cisco4948(config)#spanning-tree vlan [NUM] root secondary
PWDEMO-Cisco4948(config)#exit
PWDEMO-Cisco4948#
```

The command that actually gets entered in to Cisco running config will look like this  
spanning-tree priority 28672 but the number will be change depending on the VLAN numbers

FOUNDRY

```
BigIron(config)#spanning-tree priority 28672.
```

## **B.5 Deploy BPDU guard on all ports that use PORTFAST**

Ensure that all appropriate ports that use the PORTFAST setting are protected against BPDUs

The PORTFAST setting should only be used on ports 1G or 10G that connect clients and servers

- 10G ports which face ISIS switches should NOT use the portfast setting
- 10G ports which connect to other switches (e.g. 4900M to 4900M inter-switch link, or 4900M to cascaded 4948)

BPDU Guard will administratively shutdown any port which receives BPDUs.

```
LAB-4948-10GE-S(config)#spanning-tree portfast bpduguard default
```

This is a global command.

Also consider that when STP BPDU guard disables the port, the port remains in the disabled state unless the port is enabled manually. You can configure a port to re-enable itself automatically from the errdisable state. Issue these commands, which set the errdisable-timeout interval and enable the timeout feature:

Cisco IOS Software Commands

```
CatSwitch-IOS(config)# errdisable recovery cause bpduguard
```

```
CatSwitch-IOS(config)# errdisable recovery interval 400
```

Note: The default timeout interval is 300 seconds and, by default, the timeout feature is disabled.



For more information see:

[http://www.cisco.com/en/US/tech/tk389/tk621/technologies\\_tech\\_note09186a008009482f.shtml](http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a008009482f.shtml)

[url active MAY 2011]

Note spanning-tree bpduguard can be enabled individually on any interface which is not protected by PORTFAST

## **B5.1 Use ROOT GUARD on any interfaces that cascade to other switches**

Even though the root Primary and root secondary switch may be set, there are still situation which they can be superseded. ROOT GUARD can circumvent this.

For example in a typical 4900M HSRP configuration with cascaded 4948 switches this should be enabled on all ports 4900M ports which downlink to a 4948 or that are unused on all

spanning-tree guard root

```
LAB-4948-10GE-S(config)#interface TenGigabitEthernet1/4
LAB-4948-10GE-S(config-if)# spanning-tree guard root
LAB-4948-10GE-S(config-if)#
```

```
CISCOLAB-4948-10GE-S(config-if)#do sh run int t1/4
```

```
interface TenGigabitEthernet1/4
  switchport access vlan 10
  switchport mode access
  spanning-tree guard root
```

For more information see:

[http://www.cisco.com/en/US/tech/tk389/tk621/technologies\\_tech\\_note09186a00800ae96b.shtml](http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a00800ae96b.shtml)

[url active MAY 2011]

Using both on an interface is not required as ROOT GUARD is subordinate to BPDU GUARD

```
spanning-tree bpduguard enable
spanning-tree guard root
```

### **What Is the Difference Between STP BPDU Guard and STP Root Guard?**

BPDU guard and root guard are similar, but their impact is different. BPDU guard disables the port upon BPDU reception if PortFast is enabled on the port. The disablement effectively denies devices behind such ports from participation in STP. You must manually re-enable the port that is put into errdisable state or configure errdisable-timeout.

Root guard allows the device to participate in STP as long as the device does not try to become the root. If root guard blocks the port, subsequent recovery is automatic. Recovery occurs as soon as the offending device ceases to send superior BPDUs.



## **B.6 Use the no shutdown command on all VLANs**

Because new VLANs are shutdown [off] by default, make sure you use the `no shutdown` command.

Layer 2 interfaces are on by default

Layer 3 interfaces (routed or switched) are off by default.



PLEASE DO MAKE SURE THAT **VLAN1 IS SHUTDOWN**

## **B.7 Use the shutdown command on all unused interfaces**

Because all Layer 2 interfaces are on by default, it is a good security practice to use the `shutdown` command, and add a description as SHUTDOWN

### **CISCO**

```
PWDEMO-Cisco4948#conf t
Enter configuration commands, one per line. End with CNTL/Z.
PWDEMO-Cisco4948(config)#int g1/40
PWDEMO-Cisco4948(config-if)#desc SHUTDOWN - INTERFACE NOT USED
PWDEMO-Cisco4948(config-if)#exit
PWDEMO-Cisco4948(config)#exit
PWDEMO-Cisco4948#
```

If not being used

## **B.8 Enable secret**

It is important to apply an `enable secret` to all configs, usually during the install phase this will be a simple password such as `avid` which will be changed by the customer after system handover. Don't waste time with an `enable password` as this offers backward compatibility with Pre 12.x IOS images (last used in the 1990s!), never deployed with any Cisco Catalyst 4948 or 4900M, and it just adds confusion, also it is superseded by an `enable secret` anyway.

## **B.9 Password encryption**

Encrypting passwords is a good thing to do, but possibly this should only be done after system handover.

```
PWDEMO-Cisco4948(config)#service password-encryption
```

This will apply Cisco level 7 encryption to all clear text passwords such as line console and VTY



The `enable secret` is encrypted at Cisco level 7 by default

### **B.10 Enable telnet**

Telnet is not enabled in a Cisco 4948 by default but is enable in a Foundry by default  
A Cisco switch will need the following commands to enable telnet

```
line vty 0 4
 password avid
 login
```

A Cisco switch will need an enable secret before allowing telnet access, this is in addition to a telnet password, but the can be the same e.g. **avid**.

### **B.11 Enable synchronous logging**

When synchronous logging is enabled, information items sent to console will not interrupt the command you are typing. The command will be moved to a new line

```
line con 0
 logging synchronous
 stopbits 1
line vty 0 4
 password avid
 login
!
```

### **B.12 Get Putty 0.06**

Hyper terminal which comes with Windows is liability. It may corrupt data outside the main display window.

Putty is freeware, and v0-60 supports serial ports too. This product can be configured with a large scroll back buffer. Can be re-sized and run multiple instances to allow communications with multiple devices concurrently.

### **B.13 Logging**

- Logging is a great diagnostic tool
  - Logging is normally to a console session only
  - Logs do not persist a power cycle or reload
  - Logging can be sent to and external server
    - A syslog server
- Logging can be sent to telnet session
  - Issue the command terminal monitor
  - REMEMBER to turn it off!!
    - no terminal monitor



## B.14 Using a Syslog Server

- Many syslog implementations
  - Usually part of a commercial SNMP packages
  - Freeware applications are numerous
    - Personal favourite is KIWI Syslog Daemon
      - Which will also perform as a BASIC SNMP trap manager
        - » Just to prove they are being sent!....and received!!
  - Cisco commands

logging trap debugging THIS IS WHAT TO SEND  
 logging 10.134.87.87 THIS IS WHERE TO SEND

- Logging is only useful with time stamps
  - That means the clock must be configured
  - And the correct commands exist in the config file
  - Which is the most useful output below?

```
7w4d: %SYS-5-CONFIG_I: Configured from console by vty0 (10.134.132.86)
7w4d: %SYS-5-CONFIG_I: Configured from console by vty0 (10.134.132.86)
*Sep 14 15:39:24: %SYS-5-CONFIG_I: Configured from console by vty0 (10.134.132.86)
*Sep 14 15:39:46: %SYS-5-CONFIG_I: Configured from console by vty0 (10.134.132.86)
```

```
Aug 5 15:15:56.124 UTC: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/2, changed state to down
```

```
Aug 5 15:15:58.472 UTC: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/2, changed state to up
```

- Sep 14 at 15:39 is a whole lot more useful than 7 weeks and 4 days!!

## B.15 Timestamps

- service timestamps log uptime

```
7w4d: %SYS-5-CONFIG_I: Configured from console by vty0 (10.134.132.86)
```

- service timestamps log datetime

```
*Sep 14 15:39:24: %SYS-5-CONFIG_I: Configured from console by vty0 (10.134.132.86)
```

- Also a command for Debug statements
  - service timestamps debug uptime [datetime]
  - This parameter need to be set to if you want any sensible output when using debug commands
  - The default is uptime as no time may be set!!

```
service timestamps debug datetime
service timestamps log datetime
```

or even

```
service timestamps debug datetime msec localtime show-timezone
```



```
service timestamps log datetime msec localtime show-timezone
```

## B.16 Setting the Time

- The time can be controlled manually or by an NTP server  
Switch(config)#ntp server 10.184.106.99

Or depending on NTP setup

```
Switch(config)#ntp peer 10.184.106.99
```

Peers are a class that allows for both responses to NTP Requests as well as acceptance of NTP Updates, while an NTP Server will only respond to requests and not permitting updates.  
There is more information available from the INE's CCIEBlog -URL - <http://blog.internetworkexpert.com/tag/ntp/> - that may clear up some of the intricate detail.

- Must also set the timezone and daylight savings parameters (or just use UTC)  
clock timezone BST 1  
clock summer-time BST recurring last Sun Mar 3:00 last Sun Oct 3:00
- Manually set the clock  
Switch#clock set 17:30:00 14 SEP 2010
  - NOT an exec level command!

For Europe

```
clock timezone CET 1  
clock summer-time CET recurring last Sun Mar 2:00 last Sun Oct 2:00
```

For USA –

```
Clock summer-time EST recurring 2 Sun Mar 2:00 1 Sun Nov 2:00
```

See articles at:

<http://www.networkworld.com/community/node/11875>

And

<http://www.cisohelpcenter.com/blog/index.php/2010/03/16/daylight-savings-time-and-cisco-ios-are-you-an-hour-behind>

## B.17 Show tech support

One way to extract a lot of information form a switch is to use show tech-support

- A tool on Cisco and Foundry to do a dump of the system and all the status information
  - Shows state of all interfaces in brief and detail
    - Another reason why interface descriptions are important!
  - More information from Cisco than Foundry



- Approx 2MB from a 4948!
- Can be piped to a file on a tftp server
- 2611XM-BLUE#show tech-support | redirect tftp://10.124.87.87/redirect-sh-tech.txt

## B16.2 What is listed?

What is listed varies between switch models and s/w versions

<pre>show clock show version show running-config show stacks show interfaces show controllers show user show data-corruption show file systems show bootflash: all show cat4000_flash: all show memory statistics show process memory show process cpu show process cpu history show cdp neighbors detail show diagnostic result module all detail show environment show interfaces counters errors show interfaces status</pre>	<pre>show interfaces trunk show logging show module show mac-address-table count show platform chassis show platform cpu packet statistics show platform cpu packet driver show platform crashdump <b>show platform hardware interface all</b> VERY LONG SECTION!!! show platform health show platform environment variables show platform portmap <b>show platform software interface all</b> VERY LONG SECTION!!! show power detail show spanning-tree summary show vlan show buffers show inventory show region</pre>
--	--

## B17.2 Show tech-support - CAVEATS

- Use with caution from the Console port
  - Output of a 2MB (4948) file will be very slow and cannot be interrupted
  - It will take SEVERAL cups of coffee to finish at 9600bps!!
  - Much faster via telnet at Fast Ethernet speed
  - OK to PIPE from console port to TFTP server.
  - 2611XM-BLUE#show tech-support | redirect tftp://10.134.97.97/redirect-sh-tech.txt
  - If capturing by telnet and copy/pasting into a text file, make sure application has sufficient scroll buffer, Putty 0.6 can be configured for large amounts
  - HyperTerminal or Telnet from Window CLI ....NOT ADVISED

## B17.3 How long does it take?

- Telnet Access to switch and PIPE via Gigabit Ethernet to TFTP server

- about 10 seconds
- CLI access to a switch and PIPE via Fast Ethernet (management port) to TFTP server
  - About 60 seconds
- TELNET to CLI
  - About 3 seconds at Gigabit Ethernet or 15 seconds at Fast Ethernet
- Console to CLI
  - About 2100 seconds, yes that is 35 Minutes!

## B.18 Handover Practices

One way to ensure easy review Cisco config file documentation would be to collect the information below with three commands as part of the handover pack

```
-- show version
-- show running-config
-- show interfaces status
```

The latter command will show which interfaces are connected and which have descriptions, so connected interfaces without descriptions can be easily identified without

Below is a sample

Port	Name	Status	Vlan	Duplex	Speed	Type
Gi1/1	WEBSERVER	connected	51	a-full	a-1000	10/100/1000-TX
Gi1/2	** MC Nitris 1	notconnect	51	auto	auto	10/100/1000-TX
Gi1/3	** Protocols	connected	51	a-full	a-1000	10/100/1000-TX
Gi1/4	** MC Nitris 3	connected	51	a-full	a-1000	10/100/1000-TX
Gi1/5	** DS	connected	51	a-full	a-1000	10/100/1000-TX
Gi1/6		notconnect	51	auto	auto	10/100/1000-TX
Gi1/7	>> Interlink VLAN	connected	410	a-full	a-1000	10/100/1000-TX
Gi1/8		notconnect	51	auto	auto	10/100/1000-TX
Gi1/19		notconnect	51	auto	auto	10/100/1000-TX
Gi1/20	DOWNLINK TO 3750 V	connected	51	a-full	a-1000	10/100/1000-TX
Gi1/21		notconnect	52	auto	auto	10/100/1000-TX
Gi1/22		connected	52	a-full	a-1000	10/100/1000-TX
Gi1/23		connected	52	a-full	a-1000	10/100/1000-TX
Gi1/24		notconnect	52	auto	auto	10/100/1000-TX
Gi1/25		connected	52	a-full	a-1000	10/100/1000-TX

The additional commands are optional

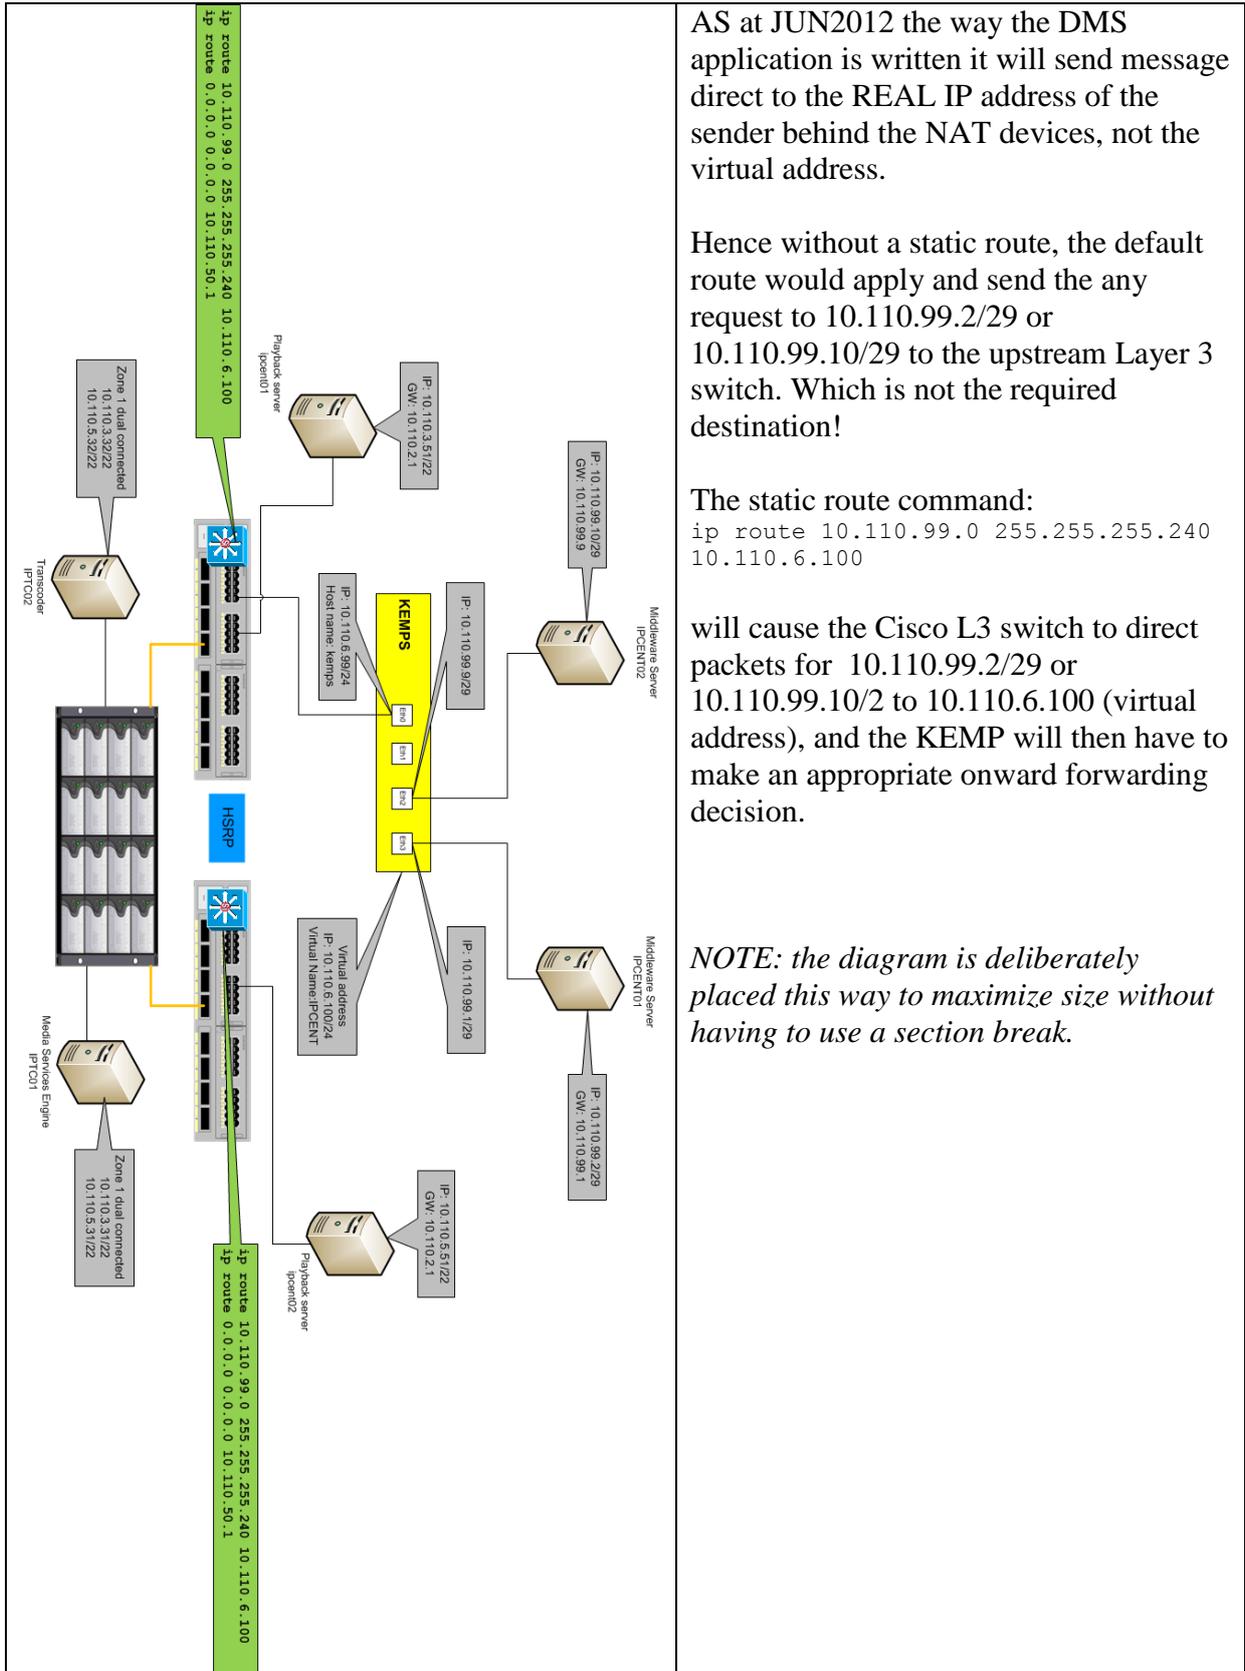
```
-- show cdp neighbors detail
-- show spanning-tree summary
-- show vlan
-- show logging
```

## B.19 Cisco Catalyst 49XX setting of the CONFIG register

Cisco 4948 switches supplied by Avid are configured with a Configuration Register value of 0x2101, which means the switch will boot from the first IOS that appears in bootflash. Cisco instructs you to set the Configuration Register to 0x2102, which means the switch will look for a boot string that points to the IOS from which to boot.



## Appendix C Interplay Central and KEMP Load Balancer



AS at JUN2012 the way the DMS application is written it will send message direct to the REAL IP address of the sender behind the NAT devices, not the virtual address.

Hence without a static route, the default route would apply and send the any request to 10.110.99.2/29 or 10.110.99.10/29 to the upstream Layer 3 switch. Which is not the required destination!

The static route command:

```
ip route 10.110.99.0 255.255.255.240 10.110.6.100
```

will cause the Cisco L3 switch to direct packets for 10.110.99.2/29 or 10.110.99.10/29 to 10.110.6.100 (virtual address), and the KEMP will then have to make an appropriate onward forwarding decision.

*NOTE: the diagram is deliberately placed this way to maximize size without having to use a section break.*

## Appendix D Fault Finding Tips – TO BE ADDED

Add here Cisco fault finding tips

## Appendix E Full Revision history

### Revision history

Note Version for this document number DOES NOT directly correlate to ISIS or Interplay Production version

Version	Name	Date	Comment
Initial Issue V1.0	David Shephard	04 July 2007	
Issue 1.1	David Shephard	01 Aug. 2007	Added section on IP requirements
Issue 1.2	David Shephard	29 Aug. 2007	Added APPX A Added Firewall info
Issue 1.3	David Shephard	15 Oct 2007	Added sections on 1.10 & 2.7 DHCP & 7.0 Security, 1.5.3 Descriptors and more on Multicast 2.2 Add section on Media converters
Issue 1.3.1	David Shephard	31 Jan 2008	Update with new URL references for Right Now Knowledge Base
Issue 1.4	David Shephard	25 June 2008	<b>Enhanced section 1.8.2 -4</b> Add information for /30 routed uplinks to CORP. <b>Added section 1.11</b> 10G link aggregation <b>Added 1.3.2</b> Add Cisco WS-X6708 <b>Added 1.3.6</b> Add Foundry Big Iron RX <b>Added 1.3.7</b> info. on inline VoIP devices <b>Added 2.3.1</b> Add info on FQDN and forward reverse lookup <b>Updated</b> section 7 links <b>Added 1.3.8</b> Discuss buffering architectures. <b>Added 5.3</b> Dual network connections. <b>Added 7.1</b> Applying Security in Network design, Mezzanine network and VMWARE. <b>Updated</b> section 1.7 Cabling with CAT 6A ratification and Valerie Rybinski article <i>NO SPECIFIC UPDATES for ISIS 1.5</i>
Issue 1.4.1	David Shephard	24 Jul 2008	<b>Section 6.3 Typo correction</b> Firewall port range for ISIS 1.4, incorrect value of 4200-4500 updated to 4200-4599
1.5.0	David Shephard	15 DEC 2008	<b>Updated section 1.9 MAN/WAN</b> <b>Update 1.10 Link aggregation</b> <b>Add section 1.12 Deploying Transfer Manager</b> <b>Add section 1.13 Jumbo Frames and Legacy applications.</b> <b>Add section 8.0 NETWORK MANAGEMENT AND MONITORING</b>

Version	Name	Date	Comment
1.6	David Shephard	02 SEP 2009	<p>Updated Section 5.3 dual connected clients</p> <p>Revise Avid URLs for new Knowledge base.</p> <p>Update information on Catalyst 4500E</p>
1.7	David Shephard	19 JUL 2010	<p>Re-order some elements in Section 1</p> <p>Insert Section 1.5 on Non Approved Switches</p> <p>Update Section 1.7 DNS</p> <p>Updated Section 1.8 Cable requirements</p> <p>Add Appendix B with Cisco switch config tips</p> <p>Add section 1.9.4 Default IP Ranges</p> <p>Add section 1.9.5 VLAN Numbering</p> <p>Add section 1.15 Avid Low Res Encoder</p> <p>Add Section 1.16 HSRP &amp; path between 4948/424</p> <p>Amendments to section 2.1 Multicast</p> <p>Update Section 2.3 DNS</p> <p>Add Section 2.8 STREAM SERVER</p> <p>Add Section 2.9 COPY SERVER</p> <p>Add Section 2.10 MOVE SERVER</p> <p>Update Section 4.0 and 4.1 Platform requirements</p>
1.8	David Shephard	04 Jan 2011	<p>Add section 1.7.1 on DNS naming</p> <p>Update section 1.9, adjust order of sections Add section 1.9.7 ISIS 5000 IP addressing</p> <p>Add Section 1.17 on INTERSWITCH Links</p> <p>Add design example 5.0.8</p> <p>Update section 6.0 Firewalling ISIS</p> <p>Add section 6.7 successfully tested firewall solutions.</p> <p>Add section 7.2 Internet connectivity restrictions.</p> <p>Add section 9: DNXHD in ZONE 4</p> <p>UPDATE APPX B Tips and good practices from V2 to V3</p>
1.9		18 AUG 2011	<p>Update sections 1.2 1.5.1, 1.3, 1.4.4, 1.4.8, 1.4.9, 1.5.7, 1.6, 1.8, 1.9.2, 2.0</p> <p>Add section 1.8.5 Patch Panels</p> <p>Add section 5.4 Using a teamed network connection</p> <p>Update top tips Appendix B add B.2.1 STP costs long/short</p> <p>Section 1.4.7 ( previously 1.5.2) Cisco Nexus 7000 approval , Other 1.4.x section incremented</p> <p>ADD 1.4.10.1 SUPER X 10G QD settings</p> <p>ADD 1.5.6.1 Foundry/Brocade MLXe</p> <p>Update 2.8.3, 2.8.4 Interplay Stream/Streaming server.</p> <p>Added section 4.4 UHRC clients</p> <p>Added section 1.6.5, 1.6.6, 1.6.7</p>
1.10			<p>Updates to appendix C</p> <p>Add new section 1.0.2 Latency impact on ISIS network traffic.</p>



Version	Name	Date	Comment
			<b>Update 1.5.2 Cisco Nexus 5500/5000/2000</b> <b>Add 1.6.9 Avid Configuration Guidelines</b> <b>Add HSRP static routing New section 1.9.6</b> <b>Add 1.18 RSTP settings for FHRP implementations.</b> <b>Update 1.9.7 Routing protocols.</b> <b>Add 1.4.12 Arista Networks 7048 – Approved switch</b> <b>Add 1.4.13 Cisco Catalyst C4500-X - Approved switch</b> <b>Add Appendix C how to configure routing for KEMP load balancer and DMS with Interplay Central Playback services</b> <b>Amend diagrams in Section 5.0.7/8</b> <b>Update section 1.4.2 Use of WS-X6708-10G-3C</b> <b>Update Avid website references with new KB URLs (where available)</b> <b>Add Section 5.0.9 and 5.0.10 with Nexus 7000 examples</b>

~END~

