

DELL EMC ISILON AND ELEMENTAL SERVER

Configuration Guide for Dell EMC Isilon Scale-Out NAS and Elemental Server v1.9

Dell EMC Solutions Group

ABSTRACT

Dell EMC Isilon and Elemental provide best-in-class, massively scalable solutions for IPTV transcoding and on-demand content delivery for broadcast media workflows. The world's top on-demand service providers rely on Dell EMC Isilon scale out NAS and Elemental Server GPU-based video transcoding solutions to simplify and accelerate their workflows.

December 2016

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2016 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA 12/16 White Paper H11206.1

Dell EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

TABLE OF CONTENTS

- EMC SOLUTIONS GROUP 1**
- ABSTRACT 1**
- EXECUTIVE SUMMARY 4**
 - Business case 4
 - Solution overview 4
 - Key results/recommendations 4
- ELEMENTAL SERVER CONFIGURATION 4**
 - Configuration 4
 - Sizing 6

EXECUTIVE SUMMARY

Business case

Elemental® and Dell EMC® Isilon® work with top service providers to convert media content for on-demand viewing on PCs, tablets, and mobile devices. Solutions from Elemental employ the massively parallel architecture of graphics processing units (GPUs) to perform accelerated video processing and encoding. Elemental reduces the number of systems needed to support the transcoding of thousands of hours of content compared to CPU-only solutions.

Service providers are faced with a number of issues as the market shifts from broadcast media content consumption to a widening array of on-demand viewing options:

- Content must be available anywhere, at any time, for any device.
- Massive libraries of legacy feature films and broadcast content require conversion.
- Valuable content must be secured and managed.

Transcoding workflows must be highly parallelized in order to prepare fresh content for on-demand delivery in a timely manner. Service providers responsible for driving new revenue from broadcast content in on-demand markets must rely on efficiencies of scale. Above all, the solutions for on-demand content creation must be cost-effective.

Solution overview

Elemental® Server is a scalable, GPU-based video processing system. Each node in an Elemental Server array is capable of converting up to 12 unique high definition or 32 mobile video files into hundreds of output files simultaneously, at data rates faster than real-time. Elemental delivers the highest quality encoding using codecs designed specifically for the GPU. Elemental Server is a Linux-based system featuring web-based administration, REST/XML APIs, and standardized SMB or NFS connectivity to high-performance Dell EMC Isilon scale-out NAS. Isilon and Elemental Server combine to provide a scalable, fault-tolerant infrastructure for media content conversion and delivery to your content delivery network.

Key results/recommendations

The world's top broadcast, streaming media, and production companies trust Isilon to store content and drive performance-hungry media workflows. Elemental and Isilon have partnered to provide a state-of-the art compute layer for on-demand transcoding workflows built on the Isilon family of scale-out NAS products. Like Isilon, Elemental Server is built on a modular appliance platform that easily scales to meet the demands of your business.

- Deploy high-density, multi-format video conversion
- Scale out network attached storage capacity and bandwidth through the simple addition of Isilon nodes
- Scale parallel GPU-based video processing power through the simple addition of Elemental Server systems

ELEMENTAL SERVER CONFIGURATION

Configuration

Elemental Server is typically deployed in a clustered configuration consisting of a management node and one or more worker nodes. The management node of the cluster acts as a gateway to balance new jobs from the REST/XML interface, watch folders, or web interface evenly across the cluster. The management node may also store an XML file to automate the configuration of new worker nodes as they are added to the Elemental Server cluster. Communication from the Elemental Server management node to the worker nodes is carried over a dedicated management network interface. The Elemental Server cluster connects to Isilon NAS over a dedicated interface on the "data network" as shown below in figure 1.

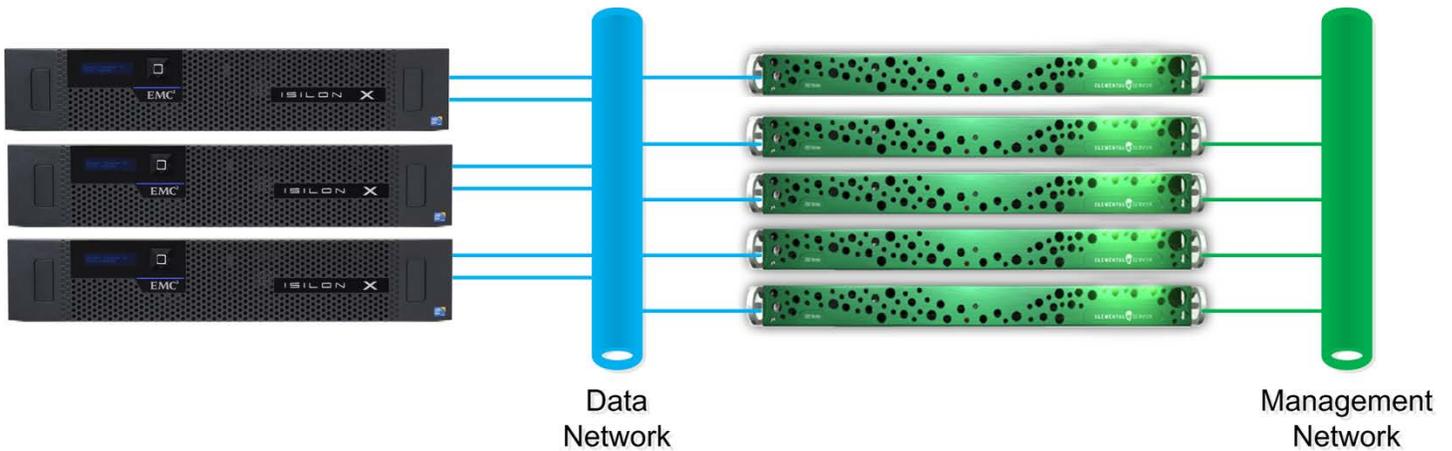


Figure 1. Elemental Server clustered network configuration

Elemental Server version 1.9 uses standard SMB and NFS mounts when connecting to Isilon NAS over the data network. Optionally, Elemental Server may connect to Isilon using DAV. DAV is only used in workflow conditions where bandwidth scalability is not critical to the workflow. The Isilon DAV protocol complies with Class 1 and 2 of the DAV specification (RFC 2518). However, all DAV clients need to go through a single node. DAV compliance is NOT met if you go through SmartConnect, or via 2 or more node IPs. DAV is an excellent protocol choice for a mixed protocol workflow, where SMB2 shares or NFS v3 exports are used as the input locations during the transcode job and a DAV is used to post a low-bitrate H.264 output file directly to a web enabled URL.

Elemental Server's current implementation of SAMBA mounts SMB shares from Isilon storage using the SMB v1 by default. SMB v1 performance when mounting Isilon is limited to about 50 Megabits per second. The transcode performance using and SMB mount on Elemental Server is roughly equivalent to the performance of a workflow utilizing the internal RAID 1 SATA storage of the Elemental Server appliance.

For best performance, Isilon storage is mounted on the Elemental Server systems using the NFS protocol via the 'mounts' section of the Elemental Server web 'settings' configuration page. When specifying the mount points to be used in transcode jobs, use the syntax `server:/ifs/export` where 'server' is the full qualified domain name of an Isilon SmartConnect zone and `/ifs/export` is the full path configured in the Isilon NFS export settings. The Isilon SmartConnect service automatically balances network connections from the Elemental Server cluster across to individual nodes in the Isilon storage cluster. For the best performance and reliability, mount the Isilon cluster from an NFS export using the SmartConnect™ Advanced application. SmartConnect Advanced allows for the dynamic allocation of IP addresses to different nodes in the NAS cluster. SmartConnect Advanced automatically initiates a failover of active NFS connections in the event of a hardware failure. SmartConnect Advanced also rebalances active NFS connections across the Isilon cluster to ensure an even distribution of active Elemental Server client connections on each node. This NFS client rebalancing is possible without disconnecting the active client connections. To learn more about SmartConnect configuration and the SmartConnect Advanced feature, please consult the Isilon [SmartConnect White Paper](#).

In addition to specifying the SmartConnect load balancing name for the Isilon cluster, specify the local mount point for the Elemental Server. Simply provide a name to `/data/mnt` prefix in the web UI and Elemental Server automatically creates the mount point. The path specified for this mount point is used when specifying media input and output locations in Elemental Server transcode job templates.

Elemental Server mounts NFS volumes via `/etc/fstab` using the following mount parameters:

```
nfs tcp,rw,async,rdirplus,rsize=32768,wsiz=32768,nfsvers=3,bg,retrans=0 0 0
```

The Elemental Server mount parameters specify that the features will be used when mounting an NFS export from Isilon storage:

- The NFS protocol uses TCP rather than UDP
- The NFS mount has read/write access
- Data communication is asynchronous
- Performance during metadata retrieval in directories with thousands of files is accelerated by combing the LOOKUP, ACCESS, GETATTR, and REaddir calls into a single REaddirPLUS call
- Data reads and writes to the NFS mount in 32 KB chunks
- The NFS v3 protocol is employed
- The mount is run in the background during startup, to prevent Elemental Server from hanging in the event the network is unavailable.

- NFS retransmission attempts are set to 0, as they are not needed when using NFS over TCP

The configuration settings for SMB shares and NFS exports require no customization when integrating Elemental Server in your workflow. When using SMB, ensure that a local user account is specified, on the cluster that matches the credentials specified for the Elemental Server SMB mount, as Elemental Server does not currently offer Microsoft Active Directory integration via the web UI. No advanced SMB parameters are required. The NFS mount parameters defined by Elemental Server are fully compatible with the default settings for a new NFS export created via the Isilon OneFS web UI.

The default Isilon “concurrency” data access pattern is well suited to the highly concurrent data access pattern of scalable transcode workflows. Please ensure that the default IO optimization settings are in place for the OneFS directory that is mounted on the Elemental Server by selecting “properties” for the subfolder from the Isilon web UI filesystem explorer page.

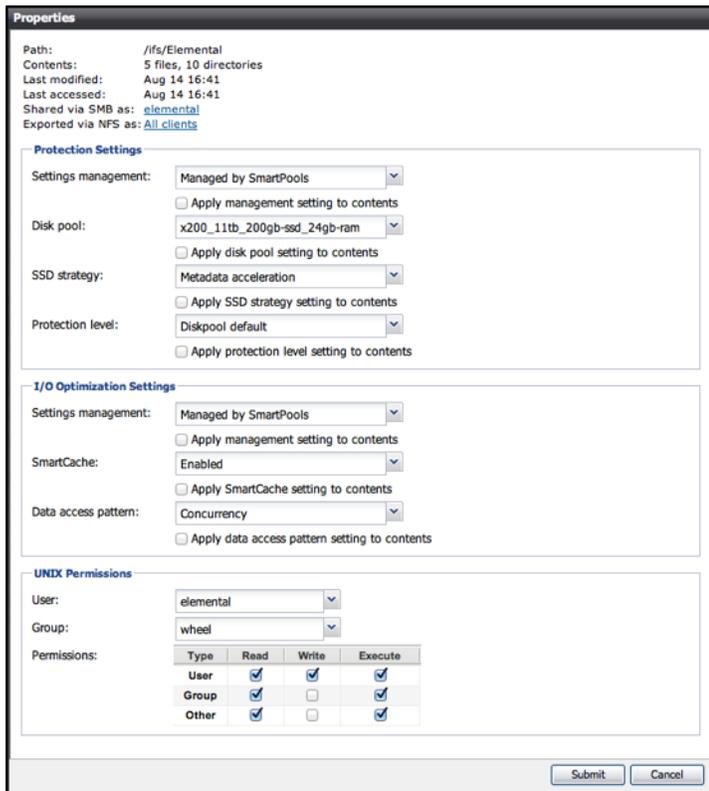


Figure 2. Default I/O optimization settings

Some throughput performance improvement is gained by changing the protection level from the Diskpool default setting of N+2:1 (the cluster can survive 2 disk failures in any node of a full node failure before going read-only) to N+1 on X200 and S200 clusters containing less than 5 nodes. Care should be taken to ensure that the protection level is increased once the cluster grows in size, as an N+1 data protection setting can extend rebuild times past an acceptable 24 hour window of safety on clusters larger than 12 nodes. Further, the N+1 data protection level should not be used on Isilon models containing 36 drive bays such as the NL and X400 series nodes.

The Isilon OneFS® operating system allows SMB shares to be mounted as NFS exports simultaneously, allowing seamless integration of Windows and UNIX/Linux applications in your workflow. Please consult the document Isilon Multiprotocol Data Access with a Unified Security Model for SMB and NFS for more information about identity mapping and simultaneous access to the same filesystem using SMB and NFS.

Sizing

The read/write throughput generated by each Elemental Server appliance is highly variable. The bandwidth requirements for Elemental Server are highly variable and “bursty” as the appliance reads files into memory.

A single 10 gigabit Ethernet enabled Elemental Server appliance with 1024 GPU cores can generate up to 3 gigabytes per second of combined read/write NFS v3 traffic when running the standard throughput benchmark test profile created by the default Elemental Server install. Using this throughput benchmark test as a general guideline, Figure 3 demonstrates example sizing ratios for an Isilon cluster against a scalable Elemental Server installation.

ISILON NODE TYPE	ELEMENTAL SERVERS PER ISILON NODE
S200 6.9 TB SAS/200GB SSD, 48 GB RAM & 10GbE	6
X400 32TB SATA/400GB SSD, 48 GB RAM & 10GbE	5
X200 11TB SATA/200GB SSD, 24 GB RAM & 10GbE	3

Table 1. Example sizing ratios

Performance characteristics of any transcode workflow will vary widely based on input codec type, output codec type, and complexity of the content video. Because of the infinite configuration and integration options available in the Isilon product family, it is important to ensure that your workflow is sized appropriately with your storage environment. Please contact your local Isilon sales representative to learn more about sizing and configuration options for the Isilon family of scale-out storage.