



DEPLOYMENT GUIDE

TriCaster® Vector AWS Deployment Guide

Welcome to the "TriCaster® Vector AWS Deployment Guide," your essential resource for deploying the TriCaster Vector solution on AWS. This comprehensive guide is crafted to streamline your journey through the landscape of AWS Customer-Deployed Edge Solutions, offering a blend of technical prowess and practical insights. Whether you are an AWS Partner seeking to enhance your customer-deployed practices or a tech enthusiast eager to dive into the world of edge solutions, this guide promises to be an invaluable companion. Inside, you'll discover detailed architectures, step-by-step deployment instructions, and strategies for optimising performance, security, and cost. Embark on a journey to elevate your deployment strategy and harness the full potential of TriCaster Vector in the AWS environment.

MARCH 2025





Contents

1	Introduction	7
1.1	Overview of TriCaster® Vectar Solution	7
1.2	Purpose of the Guide	8
1.3	Target Audience	8
2	Use Cases and Scenarios	10
2.1	Uses Cases	10
2.1.1	Live Event Broadcasting	10
2.1.2	Remote Production and Collaboration	11
2.2	Alternative Scenarios	12
2.2.1	Scaling Out an On-Prem Environment	12
2.2.2	Disaster Recovery	12
3	Solution Architecture	14
3.1	Overview of a Live Cloud Production Pipeline	14
3.1.1	The Role of TriCaster Vectar	15
3.2	Vizrt Live Cloud Production	18
3.3	Integrating TriCaster Vectar into a Live Cloud Production Pipeline	19
3.3.1	General Integration Approach for Vizrt Cloud-Ready Products	19
3.3.2	Specific Setup Requirements for TriCaster Vectar	22
3.4	High-Level Architecture Overview	23
3.5	Key Components and Services	23
3.6	Diagram of Solution Architecture	25
4	Compliance and Security	27
4.1	Data Protection and Privacy	27
4.2	User Access and Identity Management	28
4.2.1	Creating IAM Roles and Policies	29
4.3	Authenticating to your AWS EC2 instances	31



4.3.1	Secure Shell (SSH) Keys for System Administrators and DevOps	31
4.3.2	Remote desktop experience such as "NICE DCV" for Operators	31
4.3.3	TriCaster Vector LivePanel for Operators	32
4.4	Network Security	32
4.5	Increased Security Posture	34
4.5.1	Move EC2 to a Private Subnet	34
4.5.2	Implementing CIS Hardening	36
4.5.3	HTTPS for TriCaster Vector LivePanel Access	38
4.5.4	Certificate for TriCaster Vector LivePanel	39
5	Deployment and Configuration	41
5.1	Prerequisites for Deployment	42
5.1.1	Client Prerequisites	42
5.1.2	System Requirements	43
5.1.3	Skills and Knowledge	44
5.1.4	Environment Configuration	44
5.1.5	Base EC2 Instance Types and EBS Specifications	44
5.1.6	Managing AWS Service Quotas	45
5.2	Step-by-Step Manual Deployment Instructions	47
5.2.1	Deployment Instructions for Vizrt Cloud-Ready Products	47
5.2.2	Deployment Instructions for TriCaster Vector	51
5.3	Step-by-Step AWS Marketplace Deployment Instructions	53
5.4	Configuration and Customization Options	53
5.4.1	Client Network	54
5.4.2	TriCaster Vector Firewall	54
5.4.3	Firewall Ports for NDI Workflows	55
5.4.4	Firewall Ports for SRT Workflows	56
5.4.5	Firewall Ports for RTMP Workflows	56
5.4.6	Sending and receiving other protocols	57
5.4.7	Firewall for License Server	57
5.5	Alternative Deployment Options	58
5.5.1	Automated Deployment with Viz Now	58
5.5.2	Vizrt Professional Services for Cloud Migration and Support	59



6	Performance Efficiency	61
6.1	Understanding AWS Performance Metrics	61
6.2	Optimizing for Performance	62
6.3	Optimizing EC2 Instances	63
6.4	Network Performance Optimization	65
6.5	Scalability Considerations	66
6.6	Conclusion	67
7	Reliability and Resilience	69
7.1	Disaster Recovery Strategies for TriCaster Vector	71
7.2	Backup and Restore	73
7.3	Warm Standby	76
7.4	Multi-Site Active/Active	77
7.5	Monitoring and Incident Response	78
7.6	Conclusion	79
8	Operational Excellence	80
8.1	Designing for Operational Efficiency	80
8.2	Deployment Strategies for TriCaster Vector	81
8.3	Monitoring and Management	82
8.3.1	Pipeline Monitoring Hub	83
8.4	Planning EC2 GPU-based Instance Provisioning	85
8.5	Designing the Contribution Network	86
8.6	Continuous Improvement Practices	88
9	Cost Optimization	89
9.1	Budgeting and Cost Management Strategies	89
9.2	Cost-Efficient Resource Utilization	91
9.3	Understanding AWS Pricing Models	91
9.4	Storage Optimization	93
9.5	Optimizing Data Transfer Costs	94



9.6	Tools for Cost Optimization	95
9.7	AWS Calculator	96
9.8	Vizrt Flexible License Program	98
10	Integration with AWS Services	100
10.1	Integration Points with AWS Services	100
10.2	Leveraging AWS for Enhanced Capabilities	101
10.3	Example Use Cases and Scenarios	102
11	Quality Assurance and Testing	104
11.1	Testing Strategies	104
11.2	Testing and Production environments	104
11.3	Performance and Stress Testing	105
11.4	Quality Assurance Best Practices	106
12	Documentation and Support	108
12.1	User Manuals and Technical Documentation	108
12.2	Support Channels and Resources	108
13	Versioning and Updates	110
13.1	Maintaining Your Environment Secure	110
13.2	Managing Solution Updates	110
13.3	Products License	111
13.4	Emergency Maintenance	111
14	Conclusion and Next Steps	113
14.1	Summary of Key Points	113
14.2	Recommendations for Deployment Success	113
14.3	Contact Information for Further Assistance	114
15	Appendices	115
15.1	Additional Resources and References	115



15.2 Frequently Asked Questions

116



1 Introduction

1.1 Overview of TriCaster® Vector Solution

TriCaster® Vector, formerly known as Viz Vector, represents a groundbreaking evolution in live production software, uniquely designed for today's dynamic and remote-oriented video storytelling landscape. This solution, developed by Vizrt, a leader in software-defined visual storytelling tools, marks a significant shift towards virtualized live production. TriCaster Vector is a versatile, software-based platform that can be deployed on both virtual environments and on-premises, meeting a diverse range of performance requirements.

One of the most notable features of TriCaster Vector is Live Call Connect, a pioneering capability that integrates virtually every major video calling application, including Skype™, MS Teams™, Zoom Meetings™, and more. This feature significantly expands creative options for storytellers, allowing them to bring in contributors from various locations seamlessly. Additionally, the inclusion of NDI Audio I/O enables customers to virtualize their audio workflows, making TriCaster Vector a major live video production system capable of dynamic deployment in the cloud, on-premises, or in private data centres.

TriCaster Vector offers unparalleled flexibility in live production. It transcends traditional boundaries of media formats, inputs/outputs, channels, and delivery mechanisms, enabling content creation for a wide array of platforms including the Internet, mobile, and television. The system supports various aspect ratios, resolutions, and frame rates, allowing simultaneous delivery to multiple destinations. Its capabilities extend to built-in media players, recording, streaming, audio, and graphics control. Utilizing standard computer hardware and network infrastructures with IP connectivity, including SMPTE 2110 and NDI®, TriCaster Vector is adept at handling almost unlimited IP video sources.

The system's unique feature, Live Call Connect, exclusive to Vizrt Group products, elevates the production experience by enabling easy connection with in-studio guests and remote video callers, without concerns about connectivity issues. It supports a wide array of IP video sources such as SMPTE 2110, NDI®, SRT, RTMP, RTP, and more, facilitating connectivity with all types of media devices. This includes the NDI-HX Camera app, available for iPhone and Android 4K smartphones.

In essence, TriCaster Vector revolutionizes the live production domain by offering content producers unprecedented control and flexibility. It removes traditional hardware constraints on inputs, formats, locations, outputs, and modes of deployment, thus creating a new paradigm in content production that is perfectly aligned with the demands of today's multi-source, globally connected world.



1.2 Purpose of the Guide

The "TriCaster Vector Deployment Guide" is crafted to serve as a comprehensive roadmap for deploying and maximizing the capabilities of the TriCaster Vector solution within the AWS environment. This guide is designed to cater to a wide spectrum of users - from AWS partners looking to integrate TriCaster Vector into their customer-deployed edge solutions, to individual professionals aiming to leverage this powerful platform for innovative live production.

Our primary objective is to provide clear, detailed instructions and insights that simplify the deployment process, ensuring a smooth and efficient setup. We delve into the nuances of configuring TriCaster Vector to work seamlessly with AWS services, emphasizing how to harness the full potential of cloud and edge computing for live production.

Furthermore, this guide aims to enlighten users on best practices for optimizing performance, security, and cost within the AWS framework. By covering aspects like compliance, scalability, and operational excellence, we endeavor to equip readers with the knowledge needed to create robust, efficient, and secure live production environments.

In addition to the technical aspects, this guide also focuses on the strategic implications of using TriCaster Vector in various scenarios, offering insights into how this solution can transform live video production and storytelling. Whether the goal is to enhance remote production capabilities, integrate various video sources, or scale production workflows, this guide stands as a pivotal resource for achieving these objectives with TriCaster Vector on AWS.

In summary, the "TriCaster Vector Deployment Guide" is more than just a manual; it's a pathway to unlocking new possibilities in live video production, tailored for the innovative and evolving landscape of AWS-based solutions.

1.3 Target Audience

This guide is specifically tailored to meet the needs of a diverse and specialized audience, encompassing professionals and organizations at various levels of expertise and involvement in live production and cloud-based solutions. The primary target audience for this guide includes:

- **AWS Partners and Integrators:** This guide is invaluable for AWS Partners and system integrators who aim to incorporate TriCaster Vector into their customer-deployed edge solutions. It provides the necessary insights and instructions to seamlessly integrate this solution within the AWS ecosystem, enhancing their service offerings and expanding their market reach.



- **Broadcast Engineers and IT Professionals:** Professionals in broadcast engineering and IT who are responsible for setting up, managing, and optimizing live production environments will find this guide essential. It offers detailed information on the technical aspects of deploying and maintaining TriCaster Vector in alignment with AWS's well-architected framework.
- **Media Companies and Content Creators:** For media companies and independent content creators looking to leverage cutting-edge technology for live production, this guide serves as a key resource. It outlines how TriCaster Vector can be used to create dynamic, high-quality content efficiently and cost-effectively.
- **Event Producers and Directors:** Individuals involved in live event production, including producers and directors, will benefit from understanding how TriCaster Vector can enhance their production capabilities, especially in terms of integrating multiple video sources and streamlining workflows.
- **Educational Institutions and Training Centers:** This guide also caters to educational institutions and training centres that teach media production, broadcasting, and cloud technologies. It can serve as a teaching aid to help students and trainees understand the practical application of cloud-based live production tools.
- **Technical Decision-Makers:** Executives and decision-makers responsible for technology adoption and investment in organizations will find the guide useful for understanding the strategic benefits and operational implications of deploying TriCaster Vector on AWS.
- **Enthusiasts of Live Production and Cloud Technologies:** Lastly, this guide is also beneficial for enthusiasts who are keen on exploring the latest trends and technologies in live production and cloud computing.

By addressing the needs and challenges of these varied groups, this guide aims to empower a wide range of professionals and organizations to harness the capabilities of TriCaster Vector in the most effective way within the AWS cloud environment.



2 Use Cases and Scenarios

This section provides an overview of the various scenarios where TriCaster Vector adds value. It highlights the importance of these use cases in understanding the full potential of TriCaster Vector across diverse settings, showcasing its versatility and adaptability in meeting different production needs.

2.1 Uses Cases

2.1.1 Live Event Broadcasting

TriCaster Vector plays a crucial role in the broadcasting of live events such as sports, concerts, and conferences. Its integration with AWS services enhances streaming, content delivery, and scalability, providing a robust solution for managing the complexities of live event broadcasting.

Live Sports Events, Concerts, and Conferences

For live sports, music concerts, and professional conferences, TriCaster Vector facilitates the delivery of high-quality production and real-time streaming. By leveraging AWS for computing resources and content delivery networks, broadcasters can ensure reliable, high-definition broadcasts to global audiences. The platform's capacity to handle multi-source inputs and deliver real-time graphics and effects enables production teams to offer engaging and immersive live experiences.

Entertainment and Creative Productions

In TV shows, online content creation, and virtual productions, TriCaster Vector supports creativity by blending live and pre-recorded content, incorporating complex visual effects, and engaging audiences with interactive elements. Its versatility in managing diverse media formats and executing live switching and mixing positions it as a preferred tool for entertainment productions aiming to innovate beyond traditional broadcasting norms.

News and Media Broadcasting

In the fast-paced environment of newsrooms, TriCaster Vector supports the continuous and rapid broadcasting of news, handling multiple live feeds and integrating with cloud-based storage for efficient operations. Its support for geographically distributed production ensures that news organizations can maintain operations, even under challenging conditions, keeping the public informed without delays.

Key Drivers:



- **Scalability:** Utilizing AWS cloud infrastructure, broadcasters can adjust their operational scale to meet the specific demands of each live event, optimizing resource use.
- **Global Reach:** AWS's extensive network enables the delivery of live events to international audiences with reduced latency, broadening the event's audience.
- **Creative Control:** With a wide range of production tools, TriCaster Vector empowers producers to create engaging and visually appealing content.
- **Operational Resilience:** The synergy between TriCaster Vector and AWS enhances the reliability of live productions, ensuring broadcasts proceed smoothly under various conditions.

2.1.2 Remote Production and Collaboration

The solution excels in remote production setups, leveraging cloud capabilities for collaborative workflows and distributed teams. It facilitates production control from any location, offering scalability and a reduced carbon footprint. This adaptability makes it perfect for tier 1 and tier 2 broadcasters.

Key Drivers:

- **Flexibility in Location:** Moving production equipment out of traditional facilities allows for sourcing content from anywhere, with production control also being location independent.
- **Scalability:** The system can be quickly scaled up or down, offering a 'gallery on demand' approach that adapts to varying production needs.
- **Reduced Carbon Footprint:** A more sustainable production model is achieved by minimizing the need for physical transportation and reducing energy consumption in large-scale facilities.
- **Diverse Customer Base:** From enterprise customers with corporate cloud accounts to smaller, tier 2 broadcasters, TriCaster Vector serves a wide range of users.

Application:

- Outputs are streamed to end-users via CDN or returned to the continuity suite for distribution on cabled or terrestrial networks via AVoIP.



2.2 Alternative Scenarios

2.2.1 Scaling Out an On-Prem Environment

For customers embracing AVoIP and building their production facilities around this technology, TriCaster Vector can be deployed on-premises or scaled to the cloud, maintaining identical operational characteristics. It offers a cost-effective, high-quality solution for hybrid setups.

Key Drivers:

- **Hybrid AVoIP Setup:** As an AVoIP native solution, it offers a seamless operational experience, whether deployed on-premises or in the cloud.
- **Cost-Effective:** Offers a more budget-friendly solution compared to traditional SDI-based setups.
- **High Quality:** Maintains consistent video and audio quality, thanks to end-to-end NDI protocol use, avoiding quality loss from format changes or transcoding.
- **Comprehensive Solution:** Combines functionalities like video mixing, audio mixing, video server, graphics, and automation in one package.
- **Ideal for Budget-Conscious Markets:** Particularly suitable for lower tier 1 and tier 2 markets where an OPEX model is preferred.

Application

- Can be deployed on-premises on virtual machines or physical hardware and scaled to the cloud, maintaining identical operational characteristics. On-prem sources can connect to the cloud via various contribution solutions, and cloud productions can be integrated into on-prem infrastructure or distributed to end-users.

2.2.2 Disaster Recovery

TriCaster Vector stands out as an essential tool in disaster recovery scenarios. Offering an off-premises solution, it ensures continuity in production during unforeseen events or emergencies. Its deployment in AWS cloud data centers facilitates a robust and resilient setup, essential for maintaining operations during critical times.

Key Drivers:

- **Resilience in Production:** With the ability to move production equipment out of a single facility, TriCaster Vector reduces reliance on physical locations, crucial during disasters.



- **Geographic Distribution and Flexibility:** The solution supports geographically distributed production, allowing production control from any location outside of a disaster zone.
- **Scalability in Crisis:** It can be rapidly scaled up in response to disasters, adapting to the type and scale of the outage.
- **Target Audience:** Primarily aimed at enterprise customers with corporate cloud accounts.

Application:

- Outputs can be streamed to end-users via CDN or integrated into existing broadcast networks for seamless distribution, even in challenging scenarios.

Customized Solutions for Niche Markets

The solution's flexibility extends to niche markets like sports analytics, religious broadcasts, or local community programming. TriCaster Vector can be tailored to meet specific production needs and requirements, fitting lower tier 1 and tier 2 markets where budget constraints are prevalent.




3 Solution Architecture

In the evolving landscape of live production, the cloud has emerged as a pivotal platform for delivering flexible, scalable, and efficient workflows. TriCaster Vector, with its flexibility, scalability, and compatibility with cloud infrastructures, notably AWS, stands at the forefront of this revolution. Deployed as a core component of a live cloud production pipeline, TriCaster Vector enables a myriad of use cases—from broadcasting live events and managing remote productions.

This chapter delves into the architecture of live cloud production pipelines, situating TriCaster Vector within this context to illuminate its critical role. By starting with an overview of a typical live cloud production pipeline, we will progressively explore how TriCaster Vector integrates and interacts within this ecosystem. This approach will provide a foundational understanding of how TriCaster Vector contributes to achieving seamless, dynamic, and robust live production environments.

3.1 Overview of a Live Cloud Production Pipeline

A live cloud production pipeline is an intricate assembly of interconnected services and components, orchestrated to capture, produce, process, and deliver live content to global audiences. Utilizing the power of cloud technologies, this pipeline offers unparalleled flexibility in resource allocation, scalability to manage fluctuating production demands, and the agility to seamlessly integrate with a myriad of sources and distribution networks.





Content Capture: The production journey commences with the capture of live content from a diverse array of sources, including cameras, microphones, and other media inputs. This phase is crucial for gathering the raw, unprocessed feeds that will form the basis of the final product.

Content Ingestion: Following capture, the content is ingested into the cloud environment via a contribution network. This step is pivotal for transitioning the raw content from its source locations into the cloud, where it can be effectively processed and managed.

Production and Processing: Within the cloud, the raw content is transformed through a series of production processes. These processes, which include mixing, editing, and adding graphics, refine the raw feeds into polished broadcasts. It is in this crucial stage that TriCaster Vector asserts its significant role, offering a comprehensive suite of tools designed for live production in the cloud environment.

Content Distribution: After production, the content is primed for distribution. The distribution phase can encompass streaming to various online platforms, broadcasting over traditional channels, or leveraging Content Delivery Networks (CDNs) to ensure that the content reaches audiences far and wide, regardless of their location.

Management and Orchestration: Integral to the pipeline's efficacy is the use of cloud-based management and orchestration tools. These tools facilitate the smooth functioning of the pipeline, encompassing everything from resource allocation and scaling to comprehensive monitoring and security measures. A virtual control room is often instantiated, providing operators with the means to interact with pipeline components, manage signals, and route them as needed. Additionally, storage and archiving solutions are implemented to ensure both raw and produced content is preserved for future use or reference.

3.1.1 The Role of TriCaster Vector

The role of TriCaster Vector within the live cloud production pipeline is both multifaceted and pivotal, particularly during the production and processing stage. As a software-based, multi-format video and audio switcher, TriCaster Vector is tailored for contemporary production demands, facilitating remote collaboration and streamlining workflow management. Its compatibility with cloud infrastructures, especially AWS, positions it as a central tool for a broad spectrum of live production needs.

Core Features and Capabilities:

- **Live Switching, Mixing and Streaming:** TriCaster Vector stands out for its ability to perform real-time switching between multiple video source, encoding of signals (for viewing, hearing, mixing, recording and editing) with multi-bus mix effects, media playback capabilities, integrated multi viewer functionality, mix



audio feeds, incorporate graphics and streaming content to various formats, resolutions, aspect ratios and frame rates . This capability significantly enriches live content, making it more dynamic and engaging for viewers.

- **Remote Production Capabilities:** Designed with remote production in mind, TriCaster Vector supports distributed teams and sources, allowing for production control from any location. This feature is indispensable in today's globalized and mobile production environment, facilitating seamless collaboration across geographies, including scenarios of hybrid environments (Broadcaster Facilities, On Premises Remote locations and Cloud)
- **Integration with Cloud Services:** It offers smooth integration with AWS services, including Amazon EC2 for scalable compute resources, Amazon EBS, Amazon FSx and Amazon S3 for secure and durable storage, AWS Elemental MediaConnect for live video transport and AWS Elemental MediaLive for high-quality video encoding. Additionally, the solution is compatible with various AWS Network features enhancing connectivity and security within the AWS ecosystem. This ensures a cohesive workflow within the cloud production pipeline.
- **Scalability and Flexibility:** Thanks to its design, TriCaster Vector can scale its resources to match the production demand, accommodating projects of any scale, from intimate webcasts to grand broadcasts.

Advanced Specifications:

- **Versatile Input Options:** TriCaster Vector supports 44 video/audio channel inputs, accommodating a wide range of sources for comprehensive coverage and creative freedom.
- **Diverse Output Capabilities:** With 8 Mix/Effect (M/E) buses expand output versatility, allowing for the simultaneous delivery of multiple program streams. This multi-faceted capability is designed to meet intricate production demands, offering the flexibility to manage multiple HD or UHD video mix outputs.
- **Protocol Support:** It boasts compatibility with a variety of protocols including NDI®, RTMP, RTSP, HTTP and SRT, ensuring flexible source integration and content delivery.
- **Deployment Automation:** Utilizes Viz Now for automated deployment on AWS, streamlining setup and configuration processes for rapid, efficient deployment.

Technical Specifications:

- **Video Mixing:** Employs 8-M/E buses with 1 x Preview bus for advanced production planning and execution.
- **Frame Buffers:** Includes 15 frame buffers for enhanced media handling.



- **Keyable Layers:** Offers 68 keyable layers for complex compositions and effects.
- **Outputs:** Features 16 configurable outputs for versatile distribution.
- **DDR Media Players:** Comes with 4 NDI® DDR media players, supporting key and fill for enriched media playback.
- **Audio Mixer:** 16 channels of audio passthrough, integrated multi-channel audio mixer with support for 8-channel audio, DSPs and 8x8x8 audio input routing.
- **Multi-Viewers:** Support up to 4 configurable multi-viewers with workspaces and viewports for comprehensive monitoring.
- **Recording:** Supports ISO and output recording for post-production flexibility.
- **Delays:** Configurable video/audio delays at each input cater to precise timing needs.
- **Streaming Encoders:** Includes 3 integrated streaming encoders for direct live streaming.
- **Export:** Export video and image files to social media, FTP, local or external volumes, and network servers, with optional transcoding
- **Macro/Automation:** Features extensive macro and automation capabilities for streamlined operations.
- **Data Automation:** Integrated DataLink technology enabling real-time, automated data input from internal and external sources, including webpages, spreadsheets, scoreboards, databases, RSS feeds, watch files, XML, CSV, ASCII and more
- **Signal Monitoring:** Integrated waveform and vector scope, digital calibration, color preview and support for ITU-R Rec. 709
- **Tally:** Network tally via NDI
- **GPI:** Supports GPI signals (via JLCopper eBox GPI interface)
- **MIDI:** Support for standard MIDI protocol enabling third-party device control
- **Custom Panel Creation:** Allows for custom panel creation in both the operator UI and web browser, enhancing control and customization.
- **Control Options:** Offers UI control, and options for 2-stripe or 4-stripe physical control surfaces, accommodating different user preferences and operational complexities.

In the following sections, we will explore TriCaster Vectar's integration points, deployment strategies, and its utilization of AWS services to optimize live cloud production workflows further.




3.2 Vizrt Live Cloud Production

When designing their Live Cloud Production Pipeline, customers aim to select the best design that matches their business needs, often choosing the best-of-breed solutions from the market to ensure optimal performance and reliability. In this context, "Vizrt Live Cloud Production" offers a solid foundation for those embarking on this journey. It provides a carefully curated and tested selection of components, proven to work effectively across various scenarios. This offering can help customers get started with confidence, leveraging Vizrt's expertise and experience in live production to build a pipeline that meets their specific requirements.

The Vizrt Live Cloud Production Solution represents the cutting-edge of IP-based and cloud-based live production workflows, offering an end-to-end suite designed to meet the rigorous demands of modern live broadcasting.

This comprehensive bundle integrates a selection of IP-based products, encompassing video and audio mixing, replay, advanced graphics, studio automation, and intercom systems. These components are engineered to be accessible from any location, providing unparalleled flexibility and control for live production teams.

Central to this solution is Viz Now, a pivotal tool that streamlines the transition to cloud technologies for live productions. Viz Now not only facilitates the deployment of Vizrt products within the cloud but also ensures seamless integration with third-party solutions, allowing customers to craft their optimal production ecosystem with best-of-breed components.





Key Applications Include:

- **4K Video Mixing:** Offers state-of-the-art video mixing capabilities, supporting live broadcasts with high-quality, 4K video feeds, ensuring viewers receive the best visual experience.
- **Advanced Graphics and Sports Analysis:** Features broadcast-quality graphics tools that enable the creation of dynamic and engaging visual content. These tools are particularly valuable for sports broadcasts, where detailed analysis and vivid graphical representations enhance viewer engagement.
- **Studio Automation:** Provides efficient solutions for studio management and automation, enabling streamlined operations and consistent production quality with reduced manual intervention.
- **Intercom and Audio Mixing:** Incorporates high-performance audio mixing and robust intercom systems, ensuring clear and reliable communication across production teams, which is essential for coordinating complex live broadcasts.
- **Contribution and Distribution:** Offers a diverse range of solutions for moving audio and video signals from any location to any destination. This flexibility is crucial for capturing live events from around the world and distributing content to viewers across different platforms and channels.

Through its comprehensive product offerings and the continuous expansion to include the latest advancements from Vizrt and third-party solutions, the Vizrt Live Cloud Production Solution stands as a testament to the future of live broadcasting.

3.3 Integrating TriCaster Vector into a Live Cloud Production Pipeline

Integrating TriCaster Vector into a customer's live cloud production pipeline involves understanding both the general approach used for integrating Vizrt's cloud-ready products and the specific requirements and setup for TriCaster Vector. This chapter is divided into two sections to address these aspects comprehensively.

3.3.1 General Integration Approach for Vizrt Cloud-Ready Products

Network Infrastructure:

- **Virtual Private Cloud (VPC):** Provides a specialized environment within its cloud infrastructure to accommodate EC2 instances, ensuring a private segment of the AWS Cloud for your resources. This dedicated VPC, exclusive to your AWS account, guarantees isolation from other networks, creating a secure and autonomous space where you can deploy various AWS resources.




- **Subnet:** Two public subnets are specifically set up within the VPC to accommodate EC2 instances, including Tricaster Vector. These subnets are internet-facing, which simplifies the process of receiving and transmitting external content and feeds. Additionally, they facilitate seamless integration with both external services and on-premises hardware, optimizing the interaction with broader network infrastructures and services.
- **Internet Gateway:** an Internet Gateway is established as a pivotal network component within the VPC. This gateway serves as a conduit for direct online access, enabling EC2 instances to communicate with the internet. This strategic setup ensures that instances within the public subnets can send and receive data over the internet, maintaining connectivity essential for real-time content delivery and integration with various services and systems.
- **EC2 Deployment:** Within the VPC, EC2 instances are deployed in the public subnets to facilitate ease of access and initial setup. This setup serves as the foundation for deploying Vizrt cloud-ready products, including TriCaster Vector.
- **Elastic IP Allocation:** An Elastic IP address is provisioned and associated with the EC2 instance to provide a static, public IP address for remote connections as well as receive and send content / feed to external sources and destinations
- **Security Group Configuration:** The security group associated with the EC2 instance is configured to whitelist IP addresses allowed to connect to the machine, ensuring controlled access to the resources.
- **Network Access Control Lists (NACLs):** NACLs are set to allow the necessary traffic to reach the EC2 instance. The specific configurations will vary based on the needs of the production environment and will be detailed further in subsequent chapters.
- **Remote Access Connection:** For remote access connections, NICE DCV is a recommended solution. It stands out for its high-performance capabilities, which are crucial for seamless, responsive remote desktop experiences. Opting for NICE DCV ensures that users benefit from an efficient PC-over-IP (PCoIP) remote desktop functionality, which translates to a high-quality, secure, low latency and consistent user interface, regardless of network conditions or geographic distances. NICE DCV's advanced protocols deliver a rich user experience, with smooth rendering of high-definition video and audio, making it an optimal choice for demanding remote workflows.

Vizrt Licensing Management:

Licensing Options: Vizrt Licensing is based in Wibu Software Container System. Clients have two options for managing Vizrt licenses.



- The preferred involves installing a Vizrt license server within a specific VPC, referred to as the backing services VPC. This VPC should be network-peered with the VPC hosting TriCaster Vector. The method known as VPC Peering is basically a networking connection between two Virtual Private Clouds that enables you to route traffic between them using private IP addresses, essentially joining them as though they are part of the same network.
- Alternatively, clients may choose to install the license server locally within their environment, with further guidance available in the Vizrt documentation ([Vizrt Licensing Guide](#)).






3.3.2 Specific Setup Requirements for TriCaster Vector

EC2 Machine Compatibility:

- A detailed compatibility list is provided in the chapter 6 Performance Efficiency, outlining the EC2 instances that are supported by TriCaster Vector, ensuring that customers select an instance that matches their performance and resource requirements.

Elastic Block Store (EBS) Configuration:

- For optimal performance and storage capacity, it is recommended to attach at least three Amazon EBS General Purpose SSD (gp3) volumes to the EC2 instance. These volumes are designated for specific purposes:
- OS Volume: Hosts the operating system and TriCaster Vector software, ensuring fast boot and operation times. This volume is a standard gp3 (3000 IOPS and 125MB/s throughput) with minimum size of 70GB.
- Session Recording and Clips Volume: Dedicated to storing session recordings and clips, allowing for quick access and manipulation of live production content. This volume is a standard gp3 (3000 IOPS and 125MB/s throughput) with minimum size of 200GB.
- ISO Recording Volume: Utilized for ISO recordings, providing separate, dedicated storage to ensure high-quality recording of individual feeds without impacting the performance of other operations. This volume is a custom gp3 according to client needs see chapter 6 Performance Efficiency on detailed information on how to configure Storage for expected throughput and latency.






3.4 High-Level Architecture Overview

This section details a reference architecture for a small to medium Live Cloud Event Broadcasting solution leveraging Network Device Interface (NDI) for video connectivity. While this serves as a guiding framework, it's essential to recognize that multiple alternatives exist, and the architecture should ultimately be tailored to meet specific business requirements.

The solution aims to construct a comprehensive pipeline capable of multi-video and audio mixing, including a replay system and graphics capabilities, with NDI facilitating content contribution and connectivity among various components. The output is then published on third-party Content Delivery Networks (CDNs), with TriCaster Vector supporting most popular platforms including YouTube, Instagram, and others.



3.5 Key Components and Services

This list of Reference architecture components outlines a comprehensive framework for a small to medium live cloud event broadcasting solution. Each component is designed to ensure high-quality, low-latency live production through robust network infrastructure and professional-grade production tools.

1. **Content Capture:** Incorporates both NDI-native and NDI-enabled capture devices connected to cameras and microphones. These devices facilitate the conversion of analog signals, such as SDI, to NDI format, making them accessible over the network for subsequent processing and distribution.



2. **Local Network Setup:** Features a well-designed network infrastructure optimized to meet the high bandwidth demands of NDI streaming. The inclusion of high-speed switches and routers is crucial to guarantee low-latency and high-fidelity video streaming across the network.
3. **Contribution Network:** Connects various locations to the cloud, tailored to support comprehensive audio-video exchange requirements. The network's capacity must be carefully planned to handle monitoring feeds, the number of operators utilizing remote connections, and additional bandwidth for contingency purposes.
4. **Content Contribution:** Employs NDI Bridge to establish a secure link between content capture sources and the Virtual Private Cloud (VPC) designated for content processing. This setup allows for adjustable quality and latency settings to align with the bandwidth preferences specified by the customer.
5. **TriCaster Vector for Production:** Acts as the cornerstone of the live production environment, adept at managing live switching, audio mixing, and graphics integration. Its NDI compatibility ensures direct access to and seamless switching among NDI video feeds, facilitating a fluid production workflow.
6. **3Play 3PV Replay System:** An NDI-compatible replay solution seamlessly integrated into the production workflow, offering capabilities for instant replay and slow-motion playback. This system significantly enhances the production value of live events by enabling dynamic content recaps.
7. **Graphics and Branding:** Leverages the combined capabilities of Viz Trio and Viz Engine alongside TriCaster Vector to manage and overlay graphics, lower thirds, and branding elements in real-time. This integration ensures dynamic and visually engaging live broadcasts.
8. **Streaming to CDN:** The final production output from TriCaster Vector is encoded for streaming and distributed through a third-party CDN. This arrangement facilitates simultaneous streaming across multiple platforms, extending audience reach to include various social media and streaming services.

Networking Considerations:

- **Bandwidth and Quality of Service (QoS):** Adequate bandwidth provision is crucial, with QoS configurations to prioritize NDI video streaming traffic, ensuring high-quality, uninterrupted live broadcasts.
- **Security:** Implement network security practices, including secure access control and encryption where necessary, to protect the live stream content and the integrity of the production network.



3.6 Diagram of Solution Architecture

The proposed network design for a small to medium Live Cloud Event Broadcasting solution incorporates two Virtual Private Clouds (VPCs) to segregate compute resources and shared services, facilitating a streamlined and organized infrastructure.


Network Design and VPC Configuration

- **Main VPC - Live Event Broadcasting:** This VPC hosts all EC2 compute resources required for content injection, production, and management. It serves as the core environment where the live broadcasting workflow is executed, from content transcoding to signal management.
- **Second VPC - Backing Services:** Dedicated to hosting shared services like the license server and graphics archive. These services, which can be utilized across different workflows, are centralized in this VPC to provide efficiency and ease of management. A VPC peering connection is established between the Live Event Broadcasting VPC and the Backing Services VPC, facilitating secure communication between services hosted on both VPCs. The cost implications of VPC peering should be assessed based on the specific setup requirements.

Service Component	Hosted In VPC	EC2 Instance Family Type	EBS gp3 Volumes	Description
Content Transcoding NDI Bridge	Main - Live Event Broadcasting	G4dn G5	1	Facilitates efficient conversion between NDI HX compressed formats and full NDI bandwidths.
Content Processing & Video Switching TriCaster Vector	Main - Live Event Broadcasting	G4dn G5	3	Powers live switching and content management with TriCaster Vector.
Replay System Viz 3Play 3PV	Main - Live Event Broadcasting	G4dn G5	3	Enables instant replay and slow-motion playback capabilities.
Graphics Composition Viz Engine	Main - Live Event Broadcasting	G4dn G5	1	Facilitates dynamic graphics overlay and control.
Virtual Control Room & Graphics Control Viz Trio	Main - Live Event Broadcasting	G4dn G5	1	Manages graphics overlay and branding elements in real-time during broadcasts.
Signal Management - NDI Discovery Service	Main - Live Event Broadcasting	T3	1	Operates as a centralized registry of NDI sources, enhancing signal distribution management.



Vizrt License Server	Backing - Services	T3	1	Provides centralized management of product licenses for Vizrt products.
Vizrt Graphics Hub	Backing - Services	T3	2	Offers centralized management of graphics projects, ensuring efficient access to assets.





4 Compliance and Security

In today's digital age, the security and compliance of live broadcasting are paramount, especially given the high stakes surrounding broadcasting rights. For users of TriCaster Vector, prioritizing these aspects is crucial to protect content and adhere to legal and regulatory standards. This section emphasizes the importance of integrating security measures and compliance from the outset, specifically focusing on how TriCaster Vector fits into a secure live broadcasting workflow.

By highlighting TriCaster Vector's role, we aim to stress the significance of security by design and compliance in safeguarding valuable assets and maintaining trust in the competitive broadcasting industry. Adopting a proactive security posture is fundamental to the success and sustainability of live event broadcasting.

TriCaster Vector anchors its security framework on the robust principles outlined in the AWS Well-Architected Framework's Security pillar. This foundational approach ensures that TriCaster Vector's deployment and operations within AWS infrastructure inherently prioritize critical security best practices, from data protection to incident response strategies, fortifying the integrity and compliance of live broadcasting systems in the digital realm.

4.1 Data Protection and Privacy

Implementing comprehensive data protection and privacy measures is not just about fulfilling legal obligations; it's about safeguarding sensitive information and maintaining the trust of your viewers.

Key Considerations for Data Protection and Privacy:

- **Compliance with Global Regulations:** Operations in live broadcasting can cross numerous jurisdictions, each with distinct data protection laws like the General Data Protection Regulation (GDPR) in the European Union and the California Consumer Privacy Act (CCPA) in the United States. Broadcasters must be proactive in understanding and adhering to these regulations to mitigate legal risks and ensure viewer privacy. Choosing the right region for your workload based on compliance requirements is crucial for maintaining regulatory alignment.
- **Securing Sensitive Data:** The nature of broadcast content, especially content under exclusive rights, and the personal data collected from viewers necessitate rigorous security protocols. Practices such as encrypting data in transit and at rest, implementing secure access controls, and conducting regular security audits are vital to thwart unauthorized access and potential data breaches.



Vizrt Solutions and Customer Responsibilities:

- **Vizrt's Approach to Data Storage:** Vizrt solutions are designed not to store any customer-sensitive data on deployed resources. Operational data for customers are housed on EBS volumes attached to EC2 instances, ensuring a separation of sensitive information from the application layer.

Customer-Managed Encryption:

- **Data at Rest:** Customers are responsible for configuring encryption on EBS volumes to secure data at rest. Amazon EBS supports encryption across all volume types, offering transparent encryption and decryption processes that do not require manual intervention. Performance metrics, such as IOPS, remain consistent between encrypted and unencrypted volumes, with minimal impact on latency. For detailed guidance, refer to Amazon EBS encryption.
- **Data in Transit:** The encryption of data in transit is also managed by the customer. This includes ensuring that any audio or video streams are securely encrypted, with the methods varying based on the transport stream utilized. Adhering to industry best practices for the encryption of data in transit is essential for maintaining the confidentiality and integrity of broadcast content.

4.2 User Access and Identity Management

When deploying and operating TriCaster Vector within AWS, ensuring secure access to cloud resources is paramount. AWS Identity and Access Management (IAM) plays a critical role in managing access, allowing you to define who can interact with your AWS resources and how they can do so. Here's how to leverage IAM effectively during the manual deployment process, emphasizing the principle of least privilege and secure network configurations.

No Need for Root Privileges

- **Root Account Restrictions:** It's crucial to understand that deploying and operating TriCaster Vector does not require root account privileges. The use of the root account should be limited due to its unrestricted access to all resources in the AWS account. Instead, create IAM roles with specific permissions needed for deployment and management tasks. These roles defined are sets of permissions that dictate what operations can be executed within a client's AWS account. These operations are securely managed through the AssumeRole API, which enforces security by allowing only authorized entities to assume these roles, thereby minimizing the potential for unauthorized access or actions within the client's environment.

Adhering to the Principle of Least Privilege



- **Least Privilege Policy:** Follow the policy of least privilege by granting only the necessary permissions for a task and nothing more. This minimizes potential damage from accidental or malicious actions. When setting up IAM policies for users or roles involved in deploying or managing TriCaster Vector, ensure permissions are tightly scoped to the resources and actions required.
- **IAM Roles for EC2:** Utilize IAM roles to assign permissions to EC2 instances that host TriCaster Vector, enabling applications running on the instance to access AWS resources securely. This eliminates the need for storing credentials on the instance, further enhancing security.

Network Configuration and Access

- **Avoiding Public Resource Policies:** No deployment or operation of TriCaster Vector requires public resources, such as Amazon S3 buckets with policies allowing public access. Ensure all S3 buckets and other storage resources are securely configured to prevent unauthorized access, adhering to AWS security best practices.

4.2.1 Creating IAM Roles and Policies

To ensure secure and efficient deployment of your live broadcasting environment using TriCaster Vector within AWS, it's crucial to establish proper IAM roles and policies. The objective is to grant only the necessary permissions required to perform the deployment tasks, minimizing potential security risks.

Step 1: Identify Necessary Permissions

Before creating IAM roles and policies, identify the specific AWS services and resources that will be used during the deployment of TriCaster Vector and any associated components:

- Amazon EC2 instances for hosting TriCaster Vector.
- Amazon EBS for storage needs.
- Amazon VPC for network configuration.
- An optional Amazon S3 for storage of any deployment scripts or media files.

Step 2: Create an IAM Policy - TriCasterVectorDeploymentPolicy

1. **Log into the AWS Management Console** and navigate to the “Create policy” section
2. **Choose the JSON tab** and input the policy details. Use the AWS Policy Generator to create a policy that includes permissions for:
 - **EC2:** `ec2:DescribeInstances`, `ec2:StartInstances`, `ec2:StopInstances`, `ec2:RunInstances`, and any other actions necessary for managing EC2 instances.



- **EBS:** Permissions for creating, attaching, and managing EBS volumes, such as **ec2:CreateVolume**, **ec2:AttachVolume**.
- **VPC:** Necessary permissions for configuring VPCs, subnets, and security groups.
- **Optional S3:** If using S3 buckets, include permissions like **s3:ListBucket**, **s3:GetObject**, and **s3:PutObject** for the specific buckets in use.

Ensure the policy is scoped to restrict actions to only those resources necessary for the deployment tasks.

Step 3: Create an IAM Role - TriCasterVectarDeploymentRole

1. **From the IAM dashboard**, select “Roles” then “Create role.”
2. **Choose the AWS service** that will assume the role. Since this role is for deployment activities, select “EC2” as the service that will use this role, allowing EC2 instances to interact with other AWS services.
3. **Attach the Policy** you created in Step 2 to this role.
4. **Name and Review:** Give the role a descriptive name and review the permissions to ensure they align with the deployment needs.

Step 4: Assign the Role to Personnel and Resources

- For personnel, ensure they have IAM user accounts with permissions to assume the newly created role. This might involve creating an IAM policy that allows a user to assume the role and attaching it to the user or user group.
- For resources, such as EC2 instances, specify the role upon instance creation to inherit the permissions defined in the IAM role.

Best Practices

- **Regularly Review and Update:** Periodically review IAM roles and policies to ensure they remain aligned with current deployment practices and security standards.
- **Strong Password Policies:** Enforce robust password policies for IAM users to enhance account security.
- **IAM Users for Individuals:** Create individual IAM users, avoiding shared accounts for accountability.
- **Enable MFA:** Implement Multi-Factor Authentication (MFA) for enhanced security on user accounts.



- **Use IAM Groups:** For managing multiple users with similar deployment responsibilities, consider creating IAM groups and attaching policies to groups instead of individual users.
- **Enable CloudTrail Logging:** Ensure AWS CloudTrail is enabled to log all IAM actions for auditing and troubleshooting purposes.

4.3 Authenticating to your AWS EC2 instances

Ensuring secure and efficient access to AWS EC2 instances is paramount for the smooth operation of TriCaster Vector deployments. Different roles within the production team may require different methods of access, tailored to their specific responsibilities and operational needs. Here's how to approach authentication based on the operator's role:

4.3.1 Secure Shell (SSH) Keys for System Administrators and DevOps

SSH keys provide a secure method of authenticating to your AWS EC2 instances, offering a higher security level than password-based logins. This method is particularly suited for System Administrators and DevOps personnel who require direct access to the instance's command line for setup, configuration, and troubleshooting tasks.

- **Generating and Managing SSH Keys:** Utilize AWS or third-party tools to generate SSH key pairs. Store the private key securely and associate the public key with your EC2 instances upon their creation.
- **Storing SSH Key Pairs with AWS Secrets Manager:** For enhanced security and manageability, store your SSH private keys in AWS Secrets Manager. This centralizes key management and supports automated key rotation and access control through IAM policies.
- **Access Control with IAM Policies:** Define IAM policies to restrict SSH key access, ensuring only authorized personnel can retrieve and use them. Employ resource-based policies to specify access conditions.

4.3.2 Remote desktop experience such as "NICE DCV" for Operators

For operators involved in day-to-day production tasks who do not need SSH access, NICE DCV offers a secure and high-performance remote desktop experience. This allows operators to access the graphical user interface of EC2 instances running TriCaster Vector without direct instance access.

- **Configuring NICE DCV on EC2 Instances:** Ensure your instances are set up to run NICE DCV, including the installation of the necessary software and configuration of security group rules to allow access over the designated ports.



4.3.3 TriCaster Vector LivePanel for Operators

LivePanel offers a way for operators to engage with TriCaster Vector for operational tasks without requiring SSH access or direct instance interaction. Through a web-based custom user interface, LivePanel facilitates remote control of the live production process, including video mixing, media playback, audio mixing, and macro automation.

- **Accessing LivePanel:** LivePanel is accessible through web browsers, making it easy for operators to manage production tasks from any location. This web-based access simplifies the operational workflow and enhances collaboration among production team members.
- **Customization and Control:** LivePanel allows for the creation of a tailored interface that presents operators with only the tools and functions they need, streamlining the production process and reducing complexity.

Best Practices and Additional Instructions

- **Regular and Automatic Rotation of SSH Keys:** Leverage AWS Secrets Manager's automatic rotation feature to periodically update SSH key pairs, enhancing security.
- **Audit and Monitor Access:** Use AWS CloudTrail to audit access to SSH keys and NICE DCV sessions, maintaining visibility over authentication attempts and access patterns.
- **Implement Least Privilege Access:** Adhere strictly to the principle of least privilege across all access methods, ensuring users have only the permissions necessary to perform their roles effectively.

By adopting these differentiated access strategies and following best practices for security and efficiency, TriCaster Vector deployments can achieve both operational flexibility and robust security, catering to the diverse needs of the production team.

4.4 Network Security

This section explores foundational elements of network security such as Virtual Private Cloud (VPC), Network Access Control Lists (NACLs), firewalls, and intrusion detection/prevention systems (IDS/IPS), vital for safeguarding against unauthorized access and potential cyber threats.

Virtual Private Cloud (VPC)

A VPC offers an isolated segment of the AWS cloud, allowing you to launch resources within a defined virtual network. This isolation is key to controlling your network environment, including IP address range, subnet creation, and configuring route tables and network gateways. For TriCaster Vector deployments:



- **Subnet Design:** Design subnets strategically within your VPC to segregate resources based on internet exposure and architectural role. Backend systems should reside in private subnets, while public subnets can host resources like the NICE DCV Connection Gateway.
- **Public vs. Private Subnet Deployment:**
 - **Public Subnets:** Deploying EC2 instances in public subnets may be considered for better latency response. However, security measures like IP whitelisting should be employed to mitigate risks, understanding that instances will still be exposed to the internet.
 - **Private Subnets:** A more secure configuration involves placing EC2 instances in private subnets. Utilizing AWS services like NAT Gateways or AWS PrivateLink enhances security by reducing direct internet exposure while maintaining essential connectivity.
- **Security Groups:** Act as virtual firewalls at the instance level, controlling inbound and outbound traffic. Configure these groups to permit only necessary IPs and traffic for your TriCaster Vector instances.
- **VPC Internal Traffic:** To maximize efficiency and facilitate seamless integration within the VPC, the network configuration is structured to allow full communication internally. This facilitate interconnecting TriCaster Vector, the NDI Discovery Server, and any other instances present within the network. This setup is important for supporting NDI traffic.

Network Access Control Lists (NACLs)

NACLs provide an additional security layer, acting as a firewall for subnet traffic control. They are stateless, requiring separate inbound and outbound rule configurations. When setting up NACLs:

- **Stateless Nature:** Ensure rules allow legitimate traffic flow while being mindful of transport protocols like SRT and NDI, depending on your workflow needs.
- **Rule Evaluation Order:** Prioritize essential traffic through rule structuring, starting with the lowest numerical order to deny unwanted traffic explicitly.

Firewalls and IDS/IPS

Enhancing your TriCaster Vector environment's security may involve additional firewall solutions and IDS/IPS beyond AWS's built-in mechanisms:

- **AWS Network Firewall:** Offers granular traffic control, including stateful inspection and domain filtering, within your VPC.



- **Intrusion Detection and Prevention:** AWS Marketplace solutions or third-party IDS/IPS systems provide real-time monitoring for malicious activities, integrating seamlessly with AWS services.

Best Practices for Network Security

- **Least Privilege Access:** Limit network access to what's necessary for performing specific functions.
- **Regular Audits and Updates:** Continuously review and update your VPC configurations, NACLs, and firewall rules to adapt to network changes and emerging threats.
- **Encryption in Transit:** Encrypt data between your VPC and the public internet or within your VPC to prevent interception.

Increased Security Posture Model Proposal

In later sections, we will propose a secure network model tailored for live broadcasting environments. This model will focus on enhancing security without compromising on performance or accessibility, ensuring your live production workflows are protected and efficient.

4.5 Increased Security Posture

4.5.1 Move EC2 to a Private Subnet

To enhance the security posture of your live broadcasting environment while deploying TriCaster Vector on AWS, moving EC2 instances from a public subnet to a private subnet is a strategic approach. This shift significantly reduces the attack surface by limiting direct internet access to your instances.

The NICE DCV Connection Gateway serves as a bridge between clients and applications running on EC2 instances located within private subnets. By positioning the gateway within a public subnet, you enable secure client connections to the backend applications without exposing your EC2 instances directly to the internet. This setup is further complemented by the NICE DCV Session Manager, which operates on each host within your instance fleet, facilitating efficient session management and client-instance interactions.

Implementation Steps

1. **Configure EC2 Instances in a Private Subnet:**
 - Move your TriCaster Vector deployment and related EC2 instances to a private subnet within your VPC. This configuration ensures that these instances are not directly accessible from the internet, enhancing security.



2. Deploy NICE DCV Connection Gateway:

- Set up the NICE DCV Connection Gateway in a public subnet. This gateway will act as the entry point for clients, enabling them to securely access applications running on the private subnet without compromising security.
- Follow the detailed guidance provided in the AWS blog article, "[Getting started with managing NICE DCV sessions secured behind a NICE DCV Connection Gateway](#)" for step-by-step instructions on setting up the connection gateway.

3. Leverage NICE DCV Session Manager:

- Utilize the NICE DCV Session Manager, which is automatically included with your NICE DCV setup. The session manager acts as a broker, efficiently managing session requests from clients and directing them to the appropriate instances in your private subnet.

4. Security Group Configuration:

- Configure security groups for both the Connection Gateway and EC2 instances to ensure that only authorized traffic is allowed. The security group for the Connection Gateway should allow inbound connections on the NICE DCV port from clients, while outbound rules should allow connections to the instances in the private subnet.

5. Evaluate Latency Impact:

- While this approach significantly enhances security, be aware that it may introduce a minimal latency increase in the operational workflow. It's important to evaluate the impact of this latency in the context of your specific use case and operational requirements.

Benefits

- **Enhanced Security:** By isolating EC2 instances in a private subnet and using the NICE DCV Connection Gateway for access, the security of your live broadcasting environment is significantly improved.
- **Centralized Session Management:** The NICE DCV Session Manager simplifies the process of connecting clients to the appropriate instances, enhancing the efficiency of session handling.
- **Secure Remote Access:** This setup allows operators and administrators secure, remote access to the TriCaster Vector environment without direct exposure to the internet.



For detailed instructions and additional insights on configuring the NICE DCV Connection Gateway and leveraging its full potential for secure remote access, refer to the AWS blog article mentioned above. By adopting this enhanced security deployment model, customers can achieve a secure, scalable, and efficient live broadcasting setup on AWS.

4.5.2 Implementing CIS Hardening

For users aiming to enhance the security posture of their TriCaster Vector deployments on AWS, implementing Center for Internet Security (CIS) hardening on EC2 instances is a strategic move.

Understanding CIS Hardening

CIS Hardening involves configuring your EC2 instances in accordance with the CIS Benchmarks, which are consensus-based, internationally recognized security configuration guidelines. The benchmarks cover various aspects, including system and network configurations, user permissions, and software installations, aimed at minimizing vulnerabilities.

The Validation Process for TriCaster Vector

The TriCaster Vector system's compliance with CIS Benchmarks is ensured through a rigorous validation process, utilizing the assessment tool Hardening Kitty. This process involves applying CIS benchmarking guidelines and best practices for secure system configuration. Validation against CIS Benchmarks aligns TriCaster Vector with cybersecurity frameworks and regulatory standards, including NIST and the ISO 27000 series, bolstering its defenses against cyber threats.

Review and Workarounds

Below is a review of potential issues encountered during CIS application, their causes, and suggested workarounds to balance security and functionality effectively:

Inbound Connection Issues

- **Problem:** Following CIS hardening, all incoming connections are blocked.
- **Cause:** This issue arises from Windows Firewall policies within the Domain, Private, and Public profiles defaulting to block all inbound connections.
- **Workaround:** Adjusting the firewall settings to change the default behavior from "block" to "allow" for the mentioned profiles can restore essential inbound connectivity. Specifically, changing the value to "0" for these profiles helps maintain necessary security measures while ensuring connectivity.

Remote Access Challenges



- **Problem:** Post-hardening, issues with Remote Desktop and NICE DCV access emerge, particularly with default administrator credentials becoming inoperative.
- **Cause:** The root cause includes CIS benchmark settings that affect remote access and account management. Notably, certain policies may block network access for administrator accounts or suggest renaming the administrator account, impacting remote access.
- **Workaround:** To mitigate these access issues, it's advisable to comment out or adjust settings that excessively restrict remote access. Additionally, creating supplementary administrator accounts for testing and operational purposes can ensure continued access, adhering to enhanced security practices.

NDI Feed Stability

- **Problem:** Post-CIS hardening, disruptions in NDI feed stability may occur.
- **Cause:** Identified causes include specific network configuration settings and data retransmission parameters, which may adversely affect NDI feed stability.
- **Workaround:** Conducting a thorough review and subsequent modification of the CIS benchmark list—particularly adjusting settings like TcpMaxDataRetransmissions—can alleviate issues with NDI feed stability. This approach allows for the relaxation of certain network-related settings, thereby enhancing feed stability without substantially compromising on security.

Benefits of CIS Hardening

Implementing CIS Hardening for your TriCaster Vectar EC2 instances offers numerous security benefits:

- **Enhanced Security Posture:** Adhering to CIS Benchmarks significantly reduces the attack surface of your deployments, offering better protection against cyber threats.
- **Compliance:** Achieving compliance with CIS Benchmarks helps meet regulatory requirements, which is particularly important for broadcasters handling sensitive information.
- **Best Practices:** CIS Benchmarks represent the collective knowledge and expertise of cybersecurity professionals, providing a reliable foundation for securing your systems.

Applying HardeningKittyTool to TriCaster Vectar

Step 1 Download and Install HardeningKittyTool

- Download hardeningkitty tool from <https://github.com/scipag/HardeningKitty>



- Follow the instructions available in the webpage where you downloaded in the installer; section: “How To Install”
- In directory where hardeningkitty tool was installed, go to the folder “lists” and find the file **finding_list_cis_microsoft_windows_server_2022_21h2_1.0.0_machine.csv**
- Open the file found above in notepad, comment the lines mentioned below and save the file:
 - 2.2.8 (Deny access to this computer from the network)
 - 2.2.9 (Deny log on as a batch job)
 - 2.2.25 (Deny log on through Remote Desktop Services)
 - 2.2.26 (Enable computer and user accounts to be trusted for delegation)
 - 2.3.1.1 (Accounts: Block Microsoft accounts)
 - 2.3.1.4 (Accounts: Limit local account use of blank passwords to console logon only)
 - 2.3.1.5 (Accounts: Rename administrator account)
 - 2.3.7.1 ~2.3.7.9 (Remote Desktop Services settings)
 - 18.3.1 (MS Security Guide settings related to RDP)
 - 9.1.2 "Windows Firewall", "Inbound Connections (Domain Profile, Policy)"
 - 9.2.2 "Windows Firewall", "Inbound Connections (Private Profile, Policy)"
 - 9.3.2 "Windows Firewall", "Inbound Connections (Public Profile, Policy)"
- Alternatively, the file **finding_list_cis_microsoft_windows_server_2022_21h2_1.0.0_machine.csv** can be modified through the Hardening Interface. This interface provides a more user-friendly way to implement the necessary changes efficiently:
https://phi.cryptonit.fr/policies_hardening_interface/interface/windows/server/Windows%20Server%202022/21H2/Member/
- In Windows Start Menu open “Windows PowerShell”
- Run the backup command line as shown:
 - `Invoke-HardeningKitty -Mode Config -Backup -BackupFile .\<define-your-backup-file-name-here>.csv -FileFindingList .\lists\finding_list_cis_microsoft_windows_server_2022_21h2_1.0.0_machine.csv`
- Run the deployment command line as shown:
 - `Invoke-HardeningKitty -Mode HailMary -Log -Report -SkipRestorePoint -FileFindingList .\lists\finding_list_cis_microsoft_windows_server_2022_21h2_1.0.0_machine_modified.csv`
- Restart the instance.

4.5.3 HTTPS for TriCaster Vectar LivePanel Access

Given that TriCaster Vectar LivePanel is currently accessible via port 80 (HTTP), which is not encrypted and hence considered less secure, it's essential to implement measures that enhance security while awaiting the integration of HTTPS support directly within the application. The following strategy provides a temporary yet secure workaround for accessing



LivePanel through encrypted communication, without necessitating modifications to the application itself.

Employing HTTPS through External Redirection

To secure external access to the TriCaster Vector LivePanel, the system can be configured to publish TCP port 443 (HTTPS) for external connections. This method ensures that communication is encrypted, enhancing the security of data in transit. Internally, traffic received on port 443 can be redirected to TCP port 80 (HTTP), where the LivePanel operates in its current configuration. This redirection can be achieved through load balancing or reverse proxy solutions, as detailed below.

Load Balancing with AWS Elastic Load Balancer (ELB)

AWS ELB serves as an effective solution for managing this secure redirection:

- **Configuration:** Set up the ELB to listen on port 443 for incoming HTTPS traffic. Configure it to forward the traffic to the internal network on port 80, where LivePanel is hosted.
- **Benefits:** Utilizing ELB for SSL termination simplifies certificate management by centralizing it on the ELB. It also ensures that data transmitted over the internet is encrypted, providing an added layer of security.

Reverse Proxy Solutions

Alternatively, reverse proxy tools like NGINX, HAProxy, or Apache HTTP Server can be utilized for similar purposes:

- **Setup:** These servers can be configured to accept connections on HTTPS port 443 and then proxy those requests to the LivePanel running on HTTP port 80.
- **Advantages:** Reverse proxy solutions offer the flexibility of detailed request handling, SSL termination, and the capability to modify headers for enhanced security. They provide a robust framework for sophisticated traffic management and security enhancements.

4.5.4 Certificate for TriCaster Vector LivePanel

AWS Certificate Manager (ACM) is an essential service provided by AWS that simplifies the management of SSL/TLS certificates for websites and applications hosted on AWS, including the LivePanel for TriCaster Vector. SSL/TLS certificates are vital for securing network communications and authenticating the identity of websites on the internet, ensuring encrypted connections.

Integrating SSL with TriCaster Vector LivePanel



The integration capabilities of ACM with AWS services, such as AWS Elastic Load Balancing (ELB), offer significant benefits for TriCaster Vector setups. By deploying TriCaster Vector behind an AWS ELB, users can leverage ACM certificates to enable HTTPS on their LivePanel, enhancing security and efficiency in content delivery and application functionality. This integration facilitates a secure operation, automating the encryption of data in transit and simplifying the traditionally complex tasks associated with certificate management:

- **Provisioning:** ACM eliminates the need for generating certificate signing requests (CSRs), submitting them to a Certificate Authority (CA), and manual certificate installation.
- **Deployment:** Users can easily attach ACM-managed certificates to their ELB, automating the HTTPS setup for the TriCaster Vector LivePanel without the need for manual certificate installation or configuration.
- **Renewal:** ACM's automated renewal feature ensures that the SSL/TLS certificates remain up to date, providing continuous security without manual intervention.

5 Deployment and Configuration

In the dynamic landscape of live cloud event broadcasting, deploying TriCaster Vector efficiently and effectively is pivotal for leveraging its full capabilities. Recognizing the diverse needs and preferences of different organizations, TriCaster Vector supports a variety of deployment mechanisms. These range from hands-on manual approaches to more streamlined, automated solutions, catering to a spectrum of technical requirements and preferences for control versus convenience. As technology evolves, we are committed to expanding these deployment options to ensure flexibility and ease of use for our users.

This chapter delves into the various deployment models available for TriCaster Vector, offering insights into each method's unique advantages and considerations. Our goal is to provide a comprehensive understanding that empowers users to select and implement the deployment strategy that best aligns with their organizational needs and technical landscapes.

Below is an overview of the currently supported and prospective future deployment methods for TriCaster Vector, structured to facilitate a clear comparison and informed decision-making:

Deployment Model	Description	Current Support Status
Manual Deployment	Direct setup and management of resources through the AWS Management Console. Offers maximum control but demands more effort and is more prone to human error.	Supported
Automated Deployment - Viz Now	A cloud-based service facilitating the automatic deployment of software-defined workflows into AWS, optimizing live production workflows.	Supported
AWS CloudFormation	Utilizes Infrastructure as Code (IaC) to model and provision both AWS and third-party resources safely and repeatably.	Planned
Containerized Deployment	Deployment of containerized applications via Amazon ECS or Amazon EKS, allowing for scalable and efficient application management.	Not Planned
Provisioning with AMIs	Enables quick instance launching with pre-configured settings, eliminating the need for environment configuration from scratch. Contact Vizrt for arrangements.	By Agreement
AWS Marketplace	A digital catalog featuring a curated selection of software listings from independent vendors.	Supported
Vizrt Professional Services	Offers comprehensive consultancy and assistance for cloud migration, covering strategy, design, and implementation	Available



phases to facilitate a smooth transition to cloud-based production workflows.

This chapter focuses primarily on manual deployment, providing step-by-step guidance to meticulously set up TriCaster Vector within your AWS environment. This method, while more labor-intensive, grants users unparalleled control over their deployment, making it an invaluable reference for those seeking to understand and manage every aspect of their live production infrastructure. As we explore this model, we also highlight paths toward more automated and future-proof deployment strategies, ensuring your live event broadcasting capabilities remain cutting-edge and aligned with industry advancements.

5.1 Prerequisites for Deployment

Before proceeding with the deployment of TriCaster Vector and Vizrt cloud-ready products within an AWS environment, it's essential to ensure that all prerequisites and requirements are met. This section outlines the foundational elements needed for a successful deployment, including system requirements, user skills and knowledge, and environment configuration needs.

5.1.1 Client Prerequisites

To ensure the optimal performance and security of TriCaster Vector, especially when utilizing NICE DCV for remote access, certain requirements must be met on the client machines and within the network configuration. Below is a review and elaboration of these requirements:

Display Layout

- **Minimum Resolution:** A minimum display resolution of 1920x1080 is required for accessing TriCaster Vector. This ensures clarity and detail necessary for professional live production tasks.
- **Maximum Displays:** Up to 4 connected displays are supported. This limitation is set to ensure that the quality and performance of the output are not compromised, providing users with a consistent and high-quality viewing experience.

NICE DCV Client

- **Version:** The latest stable version of the NICE DCV Client is recommended. Utilizing the most recent version ensures users have access to the latest features, performance optimizations, and security enhancements, enhancing the overall remote access experience.

OS Compatibility



The following operating systems are compatible with the NICE DCV client, ensuring broad accessibility for users with different system configurations:

- **Windows 10 Pro, Enterprise, or LTSC:** These editions of Windows 10 are supported, offering a reliable foundation for the NICE DCV client.
- **Windows 11:** Full compatibility with Windows 11 allows users to leverage the latest Microsoft operating system enhancements.
- **Windows Server 2019 & Windows Server 2022:** Support for these server editions enables deployment in enterprise environments, ensuring scalability and robustness.

Security Ports – Configuration and Customization

- **Network Ports:** Access to specific network ports is a critical requirement for TriCaster Vector to function properly. The exact ports required depend on the services and features being used (e.g., remote access via NICE DCV, LivePanel access, and other networked features of TriCaster Vector).
- **Customization:** It is essential to customize your network's firewall and security settings to ensure these ports are open and accessible. This may involve configuring hardware firewalls, security software, or network policies to allow traffic on the necessary ports.
- See 5.3 Configuration and Customization Options for additional instructions.

5.1.2 System Requirements

- **Operating System:** Windows Server 2019 or 2022 is required as the operating system for the EC2 instances hosting TriCaster Vector.
- **AWS Region:** Select an AWS region that supports EBS gp3 for improved storage performance and EC2 Accelerated Computing instances, such as Amazon EC2 G4dn instances or Amazon EC2 G5 instances. These instance types offer the necessary computational power and graphics acceleration for live production workflows.
- **AWS Account:** An active AWS account is necessary to access AWS services and resources.
- **IAM User Permissions:** The AWS IAM user performing the deployment must have the necessary permissions enabled, including access to EC2, VPC, EBS, and other relevant services.
- **Available EC2 Instance Types:** Ensure the selected AWS region has the required EC2 instance types available for deployment.



5.1.3 Skills and Knowledge

- **Familiarity with AWS:** Understanding the AWS ecosystem and specific services used in the deployment is crucial.
- **AWS Networking & Content Delivery:** Knowledge in AWS Networking and Content Delivery, including VPC, Route 53, and CloudFront, is important for configuring network settings accurately.
- **AWS Compute:** Proficiency with Amazon EC2, including instance selection, configuration, and management.
- **AWS Console Proficiency:** Ability to navigate and use the AWS Console effectively for deploying and managing resources.
- **Network and IP Fundamentals:** Understanding of network, video, and audio IP fundamentals to ensure smooth setup and operation of live production workflows.

5.1.4 Environment Configuration

- **AWS Account:** Required for accessing and deploying AWS resources.
- **Vizrt License:** A valid Vizrt license is necessary for operating TriCaster Vector and other Vizrt cloud-ready products.
- **Additional Requirements:** Depending on the specific deployment scenario, additional tools or services may be required, such as network peering for VPCs or specific AWS Marketplace AMIs for third-party software integration.

5.1.5 Base EC2 Instance Types and EBS Specifications

For the initial setup of TriCaster Vector and Vizrt cloud-ready products, selecting the appropriate EC2 instance type and EBS volume specifications is crucial. These selections should be guided by the detailed recommendations provided in the Performance Efficiency chapters. This guidance will help to choose an infrastructure that aligns with the specific performance requirements of your live production workflows, considering the number and quality of signals to be managed.

- **EC2 Instance Types:** The Performance Efficiency chapters offer insights into choosing between Amazon EC2 G4dn and G5 instances. G4dn instances are suitable for scenarios with a moderate number of standard-quality signals. In contrast, G5 instances are better equipped for handling higher-quality signals or a larger quantity due to their superior computational and graphics processing capabilities. Use the information in these chapters to determine the starting configuration that best matches your production demands.



- **EBS Volume Selection:** When configuring the storage for your TriCaster Vector and Vizrt cloud-ready products deployment, selecting the right EBS volume type is critical for balancing performance with cost. As outlined in the Performance Efficiency chapters, EBS gp3 volumes should be considered the default choice due to their enhanced performance characteristics and cost efficiency. Gp3 volumes offer a significant improvement over gp2 volumes in terms of baseline performance and scalability, allowing for higher throughput and IOPS at a lower cost.
- **EBS Volume Size:** For your deployment, calculate the required storage capacity by considering both the expected hours of content to be stored and the quality of the signals. The Performance Efficiency chapters guide estimating these needs. It is essential to allocate sufficient storage to accommodate your live production workflow's demands, including considerations for future scaling.

Before finalizing your infrastructure choices, refer to the Performance Efficiency chapters to align your selections with best practices and recommendations tailored to live production environments. This preparatory step is vital to establishing a solid foundation for your deployment, ensuring it is primed for the subsequent Quality Assurance and Testing phase.

5.1.6 Managing AWS Service Quotas

An essential aspect of preparing for the deployment of TriCaster Vector and Vizrt cloud-ready products within an AWS environment is understanding and managing AWS Service Quotas. These quotas define the maximum number of resources that can be created or utilized within an AWS account or an AWS Region. To ensure your infrastructure adequately supports your live production workflow, it may be necessary to assess and, if required, request increases for specific AWS Service Quotas.

Key Service Quotas to Verify:

1. **Running On-Demand G and VT Instances:** This quota pertains to the number of virtual CPUs (vCPUs) allocated to your instances, directly impacting the computing resources available for your deployment. G4dn and G5 instances, recommended for high-performance live production environments, might require quota adjustments to meet your specific needs.
2. **EC2-VPC Elastic IPs:** Elastic IPs are crucial for maintaining static IP addresses for instances within your VPC. Verify the quota for Elastic IPs to ensure you have enough available for your deployment architecture.
3. **VPCs:** Each AWS account has a limit on the number of Virtual Private Clouds (VPCs) that can be created. Ensure this limit aligns with your network architecture needs, especially if deploying in multiple regions or requiring separate VPCs for different parts of your infrastructure.



4. **Inbound or Outbound Rules per Security Group:** Security groups control the inbound and outbound traffic to your instances. Check the limits on the number of rules per security group to ensure your security configurations can be fully implemented.
5. **Security Groups per Network Interface:** This quota affects how many security groups can be assigned to a single network interface, which could impact the networking setup of your EC2 instances.
6. **VPC Peering:** For architectures leveraging VPC peering to connect different VPCs, it's crucial to verify the quota for VPC peerings to ensure your networking architecture can be fully realized.
7. See 5.3 Configuration and Customization Options for additional instructions.

Requesting Quota Increases:

- If the default service quota limits do not meet the requirements of your live production deployment, you may need to request an increase for specific quotas.
- Navigate to the Service Quotas console within the AWS Management Console to submit these requests.
- Be aware that the approval process for quota increase requests can take up to 48 hours. Planning for potential quota increases well in advance of your deployment is advisable to avoid delays.

After selecting the appropriate EC2 instance type for your deployment, it's recommended to verify the current AWS Service Quotas against your project's requirements. Here's an example table to help with this assessment:

VPC										
Name										
Live Production										
Backing Services										
EC2 Instances										
Role	Region	VPC	VPC Multiplier	Pairing	Instance Type	Instance Family	vCPU	vCPU Increase	EIP	
License Server	EU (Frankfurt)	Backing Services	1	Live Production	t3.small	General purpose	2	0	1	
Viz Arc	EU (Frankfurt)	Live Production	1	Backing Services	m5n.2xlarge	General purpose	8	0	1	
Viz Arena	EU (Frankfurt)	Live Production	1		g5.12xlarge	GPU instance	48	48	1	
3Play	EU (Frankfurt)	Live Production	1		g4dn.2xlarge	GPU instance	8	8	1	
NDI Discovery Server	EU (Frankfurt)	Live Production	1		t3.medium	General purpose	2	0	1	
							68	56	5	
Services Quota										
VPC	Default Value	Required	Request							
VPC Peerings	50	1								
G On-demand running	0	56								
Elastic IP	5	5								



5.2 Step-by-Step Manual Deployment Instructions

When embarking on the manual installation it's important to allocate sufficient time and resources to ensure a smooth and successful setup. While the time required for each step can vary based on multiple factors, including technical complexity and personnel expertise, the total estimated time for completing the manual deployment falls within the range of 1.5 to 3 hours.

This estimate provides a comprehensive outlook, allowing for meticulous planning and execution without detailing the time allocations for individual tasks. It's designed to assist organizations in preparing for the deployment process, ensuring they allocate adequate time to address each component of the setup thoroughly and efficiently.

5.2.1 Deployment Instructions for Vizrt Cloud-Ready Products

Integrating Vizrt cloud-ready products, including TriCaster Vector, into a live cloud production pipeline requires a systematic approach. Follow these step-by-step instructions to ensure a smooth deployment within an AWS environment.

Step 1: Set Up the Virtual Private Cloud (VPC)

- Log into your AWS Management Console.
- Navigate to the VPC Dashboard and select "Create VPC."
- Define the IP address range and other configurations for your VPC. This VPC will host your cloud resources, including EC2 instances for Vizrt products.

Step 2: Deploy EC2 Instances

- Within the created VPC, navigate to the EC2 Dashboard and select "Launch Instance."
- Choose an Amazon Machine Image (AMI) that meets the system requirements for the Vizrt products you plan to deploy. It is recommended to use **Microsoft Windows Server 2019 Base** or **Microsoft Server 2022 Base** as mentioned below:


Amazon Machine Image (AMI)

Microsoft Windows Server 2019 Base
ami-035d8ae8de3734e5a (64-bit (x86))
Virtualization: hvm ENA enabled: true Root device type: ebs

Free tier eligible ▼

Description
Microsoft Windows Server 2019 with Desktop Experience Locale English AMI provided by Amazon

Architecture	AMI ID	Verified provider
64-bit (x86)	ami-035d8ae8de3734e5a	



- Creating a new *key pair*, for manual or Infrastructure as Code (IaC) deployments creating a key pair is essential for secure instance access and retrieving the Windows password via the AWS Console. Select the created key pair when launching your EC2 instance. Post-launch, use the key pair to get the initial Windows password from the EC2 console. To enhance security, store the private key file in AWS Secrets Manager following the Compliance and Security chapter's advice.
- Select an instance type from the compatibility list provided for TriCaster Vector. Ensure the instance type matches your performance and resource needs.
- Configure the instance details to deploy it within the public subnet of your VPC, enabling easier initial access and setup.

Step 3: Configure Managed prefix lists and Security Groups

- During the EC2 instance setup, define a new security group.
- After the EC2 instance launch, create a Managed Prefix List
- Add rules in the security group having the Managed Prefix List as source and specify allowed protocols and ports.
- Add whitelist IP addresses that are authorized to connect to your instance in the Managed Prefix List
- This step is crucial for ensuring secure access to your deployment.

Step 4: Allocate and Associate an Elastic IP

- Navigate to the "Elastic IPs" section within the EC2 Dashboard.
- Click on "Allocate new address" and follow the prompts to allocate an Elastic IP to your account.
- Once allocated, select the Elastic IP and choose "Associate address" to link it with your deployed EC2 instance.



Step 5: Set Up Network Access Control Lists (NACLs)

- In the VPC Dashboard, select the NACLs associated with your VPC.
- Edit the inbound and outbound rules to allow the necessary traffic for your Vizrt products to reach the EC2 instance. This configuration depends on the specific ports and protocols used by your deployment.
- See 5.3 Configuration and Customization Options for additional instructions.

Step 6: NVIDIA Driver

- Download and install the latest and stable NVIDIA Driver following the official AWS documentation. It is highly recommended to deploy GRID Drivers. More information can be found in [NVIDIA GRID Drivers](#).

Step 7: Establish a Remote Desktop Connection Using NICE DCV

- Install the latest and stable NICE DCV version on your EC2 instance following the official AWS documentation.
- Configure NICE DCV to allow remote desktop connections, ensuring you have a secure, high-performance environment for managing your Vizrt products remotely.
- **USB Driver Selection:** if you required a Control Panels and you plan to use USB remotization, ensure the USB Driver component is included in your NICE DCV installation. This component is essential for connecting on-premises Control Panels to the EC2 instance running TriCaster Vector. The USB Driver enables direct interaction with TriCaster Vector through physical control interfaces, offering an alternative or complement to the Viz Remote Connection App.
- **IDD Driver Exclusion:** For [TriCaster Vector only](#) review your NICE DCV installation to confirm that the IDD (Indirect Display Driver) is not selected or installed. While the IDD Driver offers benefits for general use by enhancing the graphics pipeline and reducing CPU usage, it conflicts with TriCaster Vector's operation. Excluding the IDD Driver is necessary to prevent potential issues where TriCaster Vector fails to open or operate correctly. This step ensures that your TriCaster Vector instance maintains its intended functionality and performance without interference from incompatible drivers.

Step 8: Create an IAM Role for NICE DCV Licensing

- **Create an IAM Policy: S3DCVLicenseBucket**
 - **Access AWS Management Console:** Log in and navigate to the IAM service, then find the "Policies" section and click "Create policy."



- **Define the Policy:**

1. Select the **JSON** tab.
2. Use the policy details provided below, or the AWS Policy Generator, to define permissions allowing **s3:GetObject** access to your specific licensing bucket. Adjust the Resource ARN to match your bucket's name and region.

```
{ "Version": "2012-10-17", "Statement": [ { "Effect": "Allow", "Action":  
"s3:GetObject", "Resource": "arn:aws:s3:::dcv-license.us-east-1/*" } ] }
```

- **Create an IAM Role: NICE DCVLICENSING**

- **Navigate to Roles:** From the IAM dashboard, select “Roles” then click “Create role.”
- **Select AWS Service:** When prompted for the type of trusted entity, choose “AWS service” and select “EC2” for the service that will assume this role. This step authorizes EC2 instances to utilize this role for interacting with other AWS services, specifically for accessing the NICE DCV licensing information.
- **Attach Policy:** Search for and attach the policy you created earlier, which includes permissions for accessing the NICE DCV licensing bucket on S3.
- **Name and Review:** Assign a descriptive name to the role, such as **NICE DCVLICENSING**, and review the configuration to ensure it meets your deployment's requirements. Confirm that the permissions align with the need to access NICE DCV licensing information securely.

Step 9: Manage Vizrt Licensing

- **Option 1 (Recommended):** Set up a Vizrt license server in a specific backing services VPC.
 - Deploy another EC2 instance in a separate VPC dedicated to licensing services.
 - Install and configure the Vizrt license server on this instance.
 - Peer the backing services VPC with the VPC hosting your TriCaster Vector and other Vizrt products to provide licensing flexibility.
- **Option 2:** Install the license server locally on the same EC2 instance as your Vizrt product, following the detailed guidance available in the Vizrt documentation.

These steps outline the general approach for deploying Vizrt cloud-ready products within an AWS environment. By following this guide, you can ensure a secure, scalable, and efficient



setup for your live cloud production pipeline, ready to integrate TriCaster Vector and other essential production tools.

5.2.2 Deployment Instructions for TriCaster Vector

Deploying TriCaster Vector within an AWS environment involves specific setup requirements to ensure optimal performance and compatibility. Follow these step-by-step instructions to configure your system accurately.

Step 1: Select the Compatible EC2 Instance and Define Elastic Block Store (EBS) Volumes

- Refer to the provided compatibility list to choose an EC2 instance type that supports TriCaster Vector. This list details the instances validated for performance and resource efficiency with TriCaster Vector.
- In **Network Settings** choose a VPC and subnet designated for your live cloud production pipeline, ensuring it is deployed in a public subnet for initial access.
- In **Configure Storage** adjust and create three volumes for the following purposes:
 - **OS Volume:** For the operating system and TriCaster Vector software. Clicking in Advanced, adjust size, IOPS and throughput following specifications mentioned on Chapter 6 of this document.
 - **Production Session and Clips:** Dedicated to storing the session (The *session* is where TriCaster stores the operating environment for the production such as video sources and format, video mixing workflow, etc...) and media clips used or produced during the broadcast. Clicking in Advanced, adjust size, IOPS and throughput following specifications mentioned on Chapter 6 of this document.
 - **ISO Recording Volume:** Used for concurrent ISO recordings, to capture individual input or output feeds. Clicking in Advanced, adjust size, IOPS and throughput following specifications mentioned on Chapter 6 of this document.
 - For all volumes mentioned above, click in Advanced, enable the column “Encrypted” and, in the column “KMS key” choose *default*.
- Launch your chosen EC2 instance within the VPC designated for your live cloud production pipeline, ensuring it is deployed in a public subnet for initial access.
- In the section “Names and Tags” it is possible to define tags for cost tracking and administration per resource types (instances, volumes and network interfaces for example). For more details consult AWS official documentation.

Step 2: Configure the EC2 Instance



- Connect to your EC2 instance using Remote Desktop (for Windows instances) to begin configuration. Detailed instructions on how connect and get windows password using AWS Console can be found in AWS documentation.
- Format and mount the attached EBS volumes to the instance, assigning each volume to its designated use case. Detailed instructions on formatting and mounting can be found in AWS documentation.
- NVIDIA Drivers: Download and install the latest and stable NVIDIA Driver following the official AWS documentation. It is highly recommended to deploy GRID Drivers. More information can be found in: [NVIDIA GRID Drivers](#)

Step 4: Install TriCaster Vectar

- Obtain the TriCaster Vectar installation package, either directly from Vizrt or through your provided licensing agreement.
- Run the installation on your EC2 instance, ensuring that TriCaster Vectar is installed on the OS volume. Follow the installation prompts to complete the setup.

Step 5: Configure Network and Security Settings

- Adjust the security group settings for your EC2 instance to open the necessary ports for TriCaster Vectar operation, including any needed for remote access, live streaming, and collaboration.
- If necessary, update the NACLs for your VPC to allow inbound and outbound traffic as required by TriCaster Vectar.
- See 5.3 Configuration and Customization Options for additional instructions.

Step 6: Licensing TriCaster Vectar

- If using a centralized Vizrt license server, ensure that your EC2 instance can communicate with the license server VPC. This may involve configuring VPC peering or setting up appropriate routing.
- Activate your TriCaster Vectar license following the instructions provided with your license, which may involve entering a license key or connecting to a license server.

Step 7: NDI Discovery Server

- If your NDI network has a NDI Discovery Server, you need to configure TriCaster Vectar to use it, before starting any test
 - On TriCaster Vectar Login Page following these steps:
 - Click in “Add-Ons”



- Navigate and click in “NDI Access Manager”
- Click in the tab “Advanced”
- In Discovery Servers/IP Address, type the IP of your NDI Discovery Server
- Close NDI Access Manager

Step 8: Final Configuration and Testing

- Once TriCaster Vector is installed and licensed, perform any final configuration steps as per your production requirements. This includes setting up live session profiles, configuring input and output settings, and customizing the environment.
- Conduct thorough testing to ensure all components are functioning correctly, including live input feeds, recording capabilities, and streaming outputs.

By meticulously following these steps, you can successfully deploy TriCaster Vector within an AWS environment, tailored to meet the demands of your live cloud production pipeline. This setup ensures you leverage the full capabilities of TriCaster Vector, providing a robust, flexible platform for your live production needs.

5.3 Step-by-Step AWS Marketplace Deployment Instructions

For those looking to deploy TriCaster® Vector through the AWS Marketplace, we've created a simplified and concise guide that provides straightforward, step-by-step instructions. The **TriCaster Vector AWS Marketplace AMI Quick Guide** is designed to help you quickly and easily launch your TriCaster® Vector instance using the AWS Marketplace.

This quick guide covers everything from initial setup and prerequisites to detailed configuration steps, ensuring a seamless deployment process. It is ideal for users who prefer clear, direct guidance that walks them through each phase of the deployment, ensuring nothing is overlooked.

The TriCaster Vector AWS Marketplace AMI Quick Guide is available in the AWS Marketplace.

5.4 Configuration and Customization Options

This section includes guidance on setting up Firewalls and Security Groups. For comprehensive insights into optimizing your EC2 instances, including choosing the right instance size and configuring EBS volumes, please refer to 6.3 Optimizing EC2 Instances.



5.4.1 Client Network

The following table outlines the necessary and optional network configuration for securely connecting from a source client network (e.g., 192.168.0.0/24) to the TriCaster Vector instance:

Application / Services	Protocol	Port	Direction	Note
NICE DCV	TCP	8443	Outbound	
NICE DCV	UDP	8443	Outbound	
TriCaster LivePanel	TCP	80	Outbound	Optional, See Compliance and Security
TriCaster LivePanel	TCP	443	Outbound	Optional, See Compliance and Security
NDI Bridge	UDP	5990	Outbound	Optional
SRT Feeds	UDP	9001-9020	Outbound	Optional
Protocols Ephemeral Ports	TCP	49152 -65535	Outbound	Optional, depends on transport protocol used
Protocols Ephemeral Ports	UDP	49152 - 65535	Outbound	Optional, depends on transport protocol used

5.4.2 TriCaster Vector Firewall

Standard firewall port configurations for TriCaster Vector.

Application Services	Protocol	Port	Direction	Note
NICE DCV	UDP	8443	Inbound	
NICE DCV	TCP	8443	Inbound	
NICE DCV	TCP	443	Inbound	
TriCaster LivePanel	TCP	80	Inbound	Optional, See Compliance and Security



TriCaster LivePanel	TCP	443	Inbound	Optional, See Compliance and Security
TriCaster API	TCP	5951	Inbound	Optional
Viz Remote Control Application	TCP	5958	Inbound	Control Surface Panel IP Connection
License Registration	TCP	22350	Inbound Outbound	See 4 Compliance and Security
Viz Flowics	TCP	80	Outbound	Optional, required when using Viz Flowics as source HTML graphics
Viz Flowics	TCP	443	Outbound	Optional, required when using Viz Flowics as source HTML graphics
rtpMIDI	TCP	5004 -5005	Inbound Outbound	Optional, supplemental audio devices as described in TriCaster Vectar use guide.
rtpMIDI	UDP	5004 -5005	Inbound Outbound	Optional, supplemental audio devices as described in TriCaster Vectar use guide.
WebRTC / TURN / STUN	TCP	3478	Outbound	Optional, for transmitting and receiving video streams without the need to expose the source or receiver directly to the internet
WebRTC / TURN / STUN	UDP	3478	Outbound	Optional, for transmitting and receiving video streams without the need to expose the source or receiver directly to the internet

5.4.3 Firewall Ports for NDI Workflows

NDI related firewall ports.

Application Services	Protocol	Port	Direction	Note
NDI Discovery Server Registration	TCP	5959	Inbound Outbound	
NDI Feed	TCP	5960 - 8999	Inbound	



Outbound				
NDI Feed	UDP	5960 - 8999	Inbound	
			Outbound	
Protocols Ephemeral Ports	TCP	49152 - 65535	Outbound	When NDI Bridge is integrated with TriCaster Vector and NDI Bridge Join Mode is set on the cloud.
Protocols Ephemeral Ports	UDP	49152 - 65535	Outbound	When NDI Bridge is integrated with TriCaster Vector and NDI Bridge Join Mode is set on the cloud

5.4.4 Firewall Ports for SRT Workflows

SRT related firewall ports.

Application Services	Protocol	Port	Direction	Note
SRT Feed	UDP	9001 – 9020	Inbound Outbound	When TriCaster Vector is used as SRT Listener
Protocols Ephemeral Ports	TCP	49152 - 65535	Outbound	When TriCaster Vector is used as SRT Caller *
Protocols Ephemeral Ports	UDP	49152 - 65535	Outbound	When TriCaster Vector is used as SRT Caller *

*Due to the nature of the SRT protocol, the ephemeral ports utilized by an SRT Caller can vary. Depending on the specific implementation of the SRT workflow, it may be necessary to use different outbound ports or enable peer-to-peer traffic. For comprehensive details, please consult the [SRT protocol RFC](#).

5.4.5 Firewall Ports for RTMP Workflows

RTMP related firewall ports.

Application Services	Protocol	Port	Direction	Note
----------------------	----------	------	-----------	------



RTMP	TCP	1935	Inbound Outbound	Used by most streaming providers such as YouTube Live, Vimeo or Twitch
-------------	-----	------	---------------------	--

5.4.6 Sending and receiving other protocols

Other protocols configurable in TriCaster Vector related firewall ports.

Application Services	Protocol	Port	Direction	Note
RTMP	TCP	80	Outbound	Periscope
RTMP	TCP	443	Outbound	Facebook or Youtube Live (when enabled under streaming configurations)
RTMP	TCP	1936	Outbound	Linkedin Live
RTMPS	TCP	2396	Outbound	Linkedin Live
RTMPS	TCP	2935	Outbound	Linkedin Live

5.4.7 Firewall for License Server

Application Services	Protocol	Port	Direction	Note
Admin Page	TCP	22352	Inbound	See 4 Compliance and Security
License Registration	TCP	22350	Inbound Outbound	See 4 Compliance and Security
NICE DCV	UDP	8443	Inbound	
NICE DCV	TCP	8443	Inbound	
NICE DCV	TCP	443	Inbound	



5.5 Alternative Deployment Options

5.5.1 Automated Deployment with Viz Now

In the realm of live production workflows, efficiency and agility are paramount. Vizrt's Viz Now platform revolutionizes the setup of live production environments by offering an automated deployment solution tailored for the dynamic needs of broadcasters and content creators. This cloud-based service streamlines the deployment process, significantly reducing the time and complexity involved in launching live production workflows.

Product Description

Viz Now is a cutting-edge, cloud-native service hosted in AWS designed to facilitate the management of live production workflows visually. It leverages pre-made templates, enabling users to swiftly set up and customize their production processes according to specific business requirements. These templates are comprehensive collections of applications configured to work in unison, allowing for extensive personalization to align with unique workflow demands.

Key Features

- **Automated Deployment:** Viz Now simplifies the deployment process, enabling the automatic setup of software-defined workflows within minutes. This automation not only saves valuable time but also reduces the potential for manual errors.
- **Customizable Applications:** Users have the flexibility to select and tailor applications within the templates to fine-tune their production workflows, ensuring that every aspect of the process aligns with their operational needs.
- **Scalable Infrastructure:** With Viz Now, adjusting storage capacity and processing power is straightforward, empowering users to optimize performance levels based on the complexity and requirements of their live production.
- **Seamless Integration and Connectivity:** The platform facilitates the integration of multiple workflows and interconnects Virtual Private Clouds (VPCs), enhancing the efficiency and flexibility of production environments.
- **Cost Efficiency:** Viz Now is offered at no additional cost to Vizrt customers who have subscribed to a Flexible Access plan, providing significant value and cost savings.

Licensing Requirements

To leverage Viz Now's capabilities, users must deploy Vizrt products that are covered under the Flexible Access licensing. Additionally, customers are required to have their own AWS (Amazon Web Services) cloud account to utilize Viz Now for deploying their live production workflows.

Time Savings and Replication



One of the most compelling advantages of Viz Now is the dramatic reduction in deployment time it offers. Completing the setup of a comprehensive Live Cloud production workflow can take as little as 15-20 minutes, a fraction of the time required for manual configurations. This efficiency not only accelerates the launch of live productions but also provides the ability to replicate the same installation programmatically across different projects or setups, ensuring consistency and reliability in production environments.

In summary, Viz Now represents a transformative approach to deploying live production workflows, marrying automation with customization to meet the evolving needs of the broadcast and content creation industries. By minimizing deployment time and enabling easy replication of setups, Viz Now positions Vizrt customers at the forefront of live production technology, ready to adapt and thrive in a fast-paced media landscape.

5.5.2 Vizrt Professional Services for Cloud Migration and Support

As the media and broadcast industry continues to shift towards cloud-based live production, Vizrt is committed to supporting our customers through every step of their cloud journey. With the Vizrt Professional Services for Cloud Migration and Support, we offer an initiative designed to leverage our extensive expertise in cloud technologies and live production workflows. This program draws on our successful experiences to provide customers with the guidance and support needed for a seamless transition to the cloud.

Description

Our Professional Services for Cloud Migration and Support are tailored to deliver end-to-end consultancy and assistance throughout the strategy, design, and implementation phases of cloud migration. Recognizing the complexities of migrating live production environments to the cloud, our dedicated team of experts is here to navigate our customers through this transformative process.

Strategy and Design Consultancy

The journey begins with an in-depth consultancy phase, where our specialists collaborate with customers to understand their specific needs, business objectives, and technical environment. By evaluating the current infrastructure and identifying opportunities for cloud migration, we offer strategic advice on utilizing cloud technologies to meet business goals. This collaborative approach ensures the development of a customized cloud migration strategy and the design of a scalable architecture that precisely fits the customer's requirements.

Implementation Support

Following the strategy and design phase, our professional services team provides hands-on support for the implementation of the migration plan. Our experienced cloud engineers and architects work closely with the customer's IT team to facilitate a smooth transition to the



cloud. From setting up the necessary infrastructure and configuring cloud services to migrating applications and data, we ensure every aspect of the cloud environment is optimized for live production. Our team offers best practices, security measures, and cost optimization strategies to enhance the cloud migration experience.

Benefits

- **Expert Guidance:** Benefit from the deep expertise of our consultants and cloud engineers who bring a wealth of knowledge in cloud technologies and live production workflows.
- **Tailored Solutions:** Receive customized strategies and architectures developed to meet the unique needs and challenges of your organization.
- **Reduced Risk:** Minimize the risks associated with cloud migration through our structured and expert-led approach, ensuring a successful transition.
- **Accelerated Time-to-Value:** Experience a streamlined migration process that allows for quicker realization of cloud benefits, enhancing your operational capabilities.
- **Cost Optimization:** Leverage our insights into cost-effective cloud strategies to achieve operational efficiencies and maximize your investment.

6 Performance Efficiency

Achieving performance efficiency is crucial for broadcasters using TriCaster Vector on AWS, as it ensures smooth and high-quality live broadcasts. This core objective involves optimizing cloud resources for real-time video processing and streaming, balancing power with cost-effectiveness and scalability. Key strategies include selecting appropriate AWS resources, configuring them optimally, and adjusting to demand.

Performance efficiency also means managing resources to avoid over or under-provisioning during live events. It hinges on understanding AWS metrics, conducting regular tests, and monitoring to swiftly tackle bottlenecks. This guide aims to equip broadcasters with strategies to enhance their TriCaster Vector deployments, emphasizing resource optimization, network tuning, and scalable solutions for cost-efficient, high-quality live production.

6.1 Understanding AWS Performance Metrics

In the realm of live cloud production closely monitoring performance metrics is essential for maintaining optimal operational efficiency. AWS provides a suite of metrics and tools designed to help users monitor and analyze deployments' performance in real-time. Understanding these metrics and effectively utilizing AWS's monitoring tools can significantly enhance the performance and reliability of live broadcasts.

Key AWS Metrics for Monitoring Performance

1. **CPU Utilization:** Measures the percentage of allocated compute units that are currently in use by an instance. High CPU utilization may indicate that an instance is overworked and might require scaling.
2. **Network In/Out:** Tracks the amount of data moving into and out of your AWS resources. Monitoring network throughput is crucial for live production workflows to ensure there are no bottlenecks that could impact stream quality.
3. **Disk I/O Operations:** Provides insights into the read and write operations on the disk, important for understanding the data processing and storage performance of your instances.
4. **Latency:** Measures the time taken for a data packet to travel from source to destination. Low latency is critical for live streaming to ensure real-time performance.
5. **Error Rates:** Tracks the number of errors encountered during the operation of your AWS resources. High error rates can indicate underlying issues with your setup or configuration.



Tools for Performance Monitoring and Analysis

- **Amazon CloudWatch:** A comprehensive monitoring service that provides detailed insights into AWS resource utilization, application performance, and operational health. CloudWatch allows you to collect and track metrics, set alarms, and automatically react to changes in your AWS resources.
- **AWS CloudTrail:** Helps in governance, compliance, operational auditing, and risk auditing of your AWS account. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.
- **AWS Trusted Advisor:** An online tool that offers real-time guidance to help you provision your resources following AWS best practices. Trusted Advisor inspects your environment and makes recommendations to help improve performance by optimizing AWS infrastructure, improving security and compliance, reducing costs, and monitoring service limits.

6.2 Optimizing for Performance

For broadcasters utilizing TriCaster Vector on AWS, achieving optimal performance and efficiency is paramount for the success of live cloud productions. This necessitates a strategic approach to selecting and configuring EC2 instances, ensuring they are perfectly tailored to meet the demands of live broadcasting workflows.

Below are key considerations for optimizing your AWS deployment for TriCaster Vector.

- **Content Ingestion and Distribution Design:** While TriCaster Vector comes with integrated streaming decoders and encoders, larger deployments might benefit from decoupling these stages. This approach offers enhanced flexibility and allows for a more balanced load distribution across various components.
- **Signal Management:** Evaluate the number of input and output signals your production will handle. This assessment is crucial for selecting instances that can manage the anticipated load without compromising performance.
- **Transport Protocol, Signal Quality, and Processing Needs:** The choice of transport protocol, along with signal quality, directly influences whether CPU or GPU resources—or a combination of both—are necessary. Accurately estimating the maximum encoding and decoding load is vital for ensuring your system can handle peak demands efficiently.
- **Functional Requirements:** Define the specific functionalities you plan to utilize, such as content recording and media playback. Each feature has its own set of resource requirements that must be considered when choosing your instances.



- **Remote desktop Requirements:** The requirements for remote desktop services should also be taken into account. This includes understanding the bandwidth, latency, and computational resources needed to ensure a seamless remote user experience.

The following recommendations are based on tested scenarios and loads, offering a solid starting point for broadcasters to configure their environments effectively. These settings are designed to meet a broad range of live production demands, ensuring smooth operation and high-quality content delivery.

6.3 Optimizing EC2 Instances

Recommended infrastructure and settings for leveraging AWS accelerated computing and GPU instances, alongside best practices for configuring instances and storage solutions for optimal performance.

Accelerated Computing Instances

- **Recommended Instances:** For compute-intensive tasks such as video processing, Amazon EC2 G4dn and EC2 G5 instances have been tested and selected for their high compute efficiency. These instances are equipped with NVIDIA GPUs, offering substantial acceleration for video processing tasks.
- **Stay Informed with AWS Instance Types Page:** AWS regularly updates the descriptions and use cases for each instance type. Refer to this resource to align your workload requirements with the capabilities of different instance types, ensuring you select the most appropriate based on tested instances.

Utilizing GPU Instances for Enhanced Video Processing

- **GPU Requirements Identification:** TriCaster Vector's video processing and encoding for live productions benefit significantly from GPU acceleration. Currently, it supports a single NVIDIA GPU, with plans to enable multiple GPU support in the future.
- **Selecting GPU Instances:** Ensure the chosen GPU instances are properly configured for your workload. This entails installing the necessary drivers and libraries, with details provided in the Deployment and Configuration section for the recommended driver version.

Configuring Instances for Optimal Performance

- **Enhanced Networking:** Activate Enhanced Networking with an Elastic Network Adapter (ENA) to reduce latency, jitter, and CPU utilization, critical for ensuring smooth video processing and streaming.



- **EBS Optimized Instances:** Utilize EBS-optimized instances to enhance block storage performance. These instances offer dedicated bandwidth to EBS volumes, improving I/O performance significantly.

Selecting the Right Storage Solutions

- **Amazon EBS gp3 Volumes:** Recommended for high-performance applications requiring frequent read/write operations with low latency. gp3 volumes are ideal for storing media assets, providing consistent and fast access.

Configuring Storage for High Throughput and Low Latency

- **General Purpose SSD (gp3) Volumes:** Suggested configurations for gp3 volumes are detailed in the accompanying table, tailored to various storage purposes. Adjust IOPS and throughput settings based on specific needs.
- **Storage Capacity Planning:** Estimate the required storage capacity by considering the expected hours of content and the signal quality. The Performance Efficiency chapters offer guidance on these estimations, ensuring adequate storage is allocated for your live production workflow, with scalability in mind.

Vizrt-Tested EC2 Instances

Instance	GPUs	vCPU	Memory (GiB)	GPU Memory (GiB)	Instance Storage (GB)	Network Performance (Gbps)	EBS Bandwidth (Gbps)
g4dn.2xlarge	1	8	32	16	1 x 225 NVMe SSD	Up to 25	Up to 3.5
g4dn.4xlarge	1	16	64	16	1 x 225 NVMe SSD	Up to 25	4.75
g4dn.8xlarge	1	32	128	16	1 x 900 NVMe SSD	50	9.5
g5.2xlarge	1	24	8	32	1 x 450 NVMe SSD	Up to 10	Up to 3.5
g5.4xlarge	1	24	16	64	1 x 600 NVMe SSD	Up to 25	8
g5.8xlarge	1	24	32	128	1 x 900 NVMe SSD	25	16
g6.2xlarge	1	8	32	24	1x450 NVMe SSD	Up to 10	Up to 5
g6.4xlarge	1	16	64	24	1x600 NVMe SSD	Up to 25	8
g6.8xlarge	1	32	128	24	2x450 NVMe SSD	25	16



Optimal EC2 Instance Selection by Use Case

Instance	Inputs	Output	NICE DCV Monitors	Functionality
g4dn.2xlarge	4 x NDI HQ	8 X NDI HQ 1 x SRT	2	Recording on NDI Input and Outputs Only
g4dn.4xlarge	8 x NDI HQ 2 x SRT	8 X NDI HQ 2 x SRT	3	Recording on NDI and SRT inputs NDI Outputs Format conversion SRT Output Encoding
g4dn.8xlarge	9 x NDI HQ 2 x SRT	8 x NDI HQ 2 x SRT 1 x <i>other protocols</i>	4	Recording on NDI and SRT inputs NDI Outputs Format conversion SRT Output Encoding Streaming other protocols

NDI HQ = 105 -132 Mbps (1920 x 1080p60)

NiceDCV = 10 Mbps – 21Mbps per monitor (1920 x 1080p60)

SRT = 6.9 Mbps (1920 x 1080p60)

Recommended EBS Volume configurations

EBS volume types	Purpose	Size GB	IOPS	Throughput MiB/s	Recording Time	Scenario
gp3	OS and SW	70	3000	125		
gp3	Sessions and Clips	200	3000	125		
gp3	Recordings	500	3000	150	1h 30min	6 recordings HD (1920x1080p60)
gp3	Recordings	500	3000	200	1h	7 recordings HD (1920x1080p60)
gp3	Recordings	500	3200	400	30min	6-8 recordings 4K (3840x2160p60)

6.4 Network Performance Optimization

A well-designed network architecture, effective bandwidth management, and leveraging AWS services are key components of achieving optimal network performance. Here's a guide to enhancing your network for TriCaster Vector deployments.



Designing High-Performance Network Architectures

- **Segment and Structure:** Design your network to segment different types of traffic (e.g., live feed, monitoring, management) using Amazon Virtual Private Cloud (VPC). This allows for more precise control over traffic flow and security.
- **Direct Connectivity:** For consistent high performance, consider using AWS Direct Connect. It establishes a dedicated network connection from your premises to AWS, reducing latency, increasing bandwidth, and providing a more consistent network experience than internet-based connections.

Bandwidth Management and Optimization Strategies

- **Content Delivery Network (CDN):** Utilize Amazon CloudFront to distribute your live streams. A CDN minimizes latency by caching content closer to your viewers, efficiently managing bandwidth by offloading the traffic from your origin server.
- **Adaptive Bitrate Streaming:** Implement adaptive bitrate streaming to dynamically adjust the quality of a video stream in real time according to a user's network speed. This approach helps in optimizing bandwidth usage and improving viewer experience.
- **Monitor and Allocate Bandwidth:** Regularly monitor your network usage with Amazon CloudWatch to identify patterns and peak usage times. Allocate bandwidth accordingly to prioritize critical traffic, especially during live events.

Leveraging AWS Services for Improved Network Efficiency

- **Network Access Control Lists (NACLs) and Security Groups:** Configure NACLs and Security Groups to control inbound and outbound traffic at both the subnet and instance level. Proper configuration can help in safeguarding against unwanted traffic, reducing the risk of network congestion.

6.5 Scalability Considerations

The TriCaster Vector scalability focus is on strategic monitoring, adjustment of resources, and architectural decisions that enhance performance and reliability. Here's how to approach it.

Continuous Performance Monitoring and Adjustment

- **Performance Metrics:** Implement a robust monitoring system using tools like Amazon CloudWatch to continuously track performance metrics such as CPU utilization, network bandwidth, and storage I/O operations. These metrics provide valuable insights into the system's current load and performance.
- **Resource Adjustment:** Based on the monitored metrics, adjust computing power and storage resources proactively to meet the demands of your live production. This



might involve scaling up EC2 instances or provisioning additional EBS volumes before large-scale events to ensure smooth operation.

Decoupling Decoding and Encoding Services

- **Service Decoupling:** In larger productions, it's beneficial to decouple decoding and encoding tasks from the main TriCaster Vector pipeline. Managing these tasks as separate services allows for more flexible resource allocation and easier scaling based on specific needs.
- **Leveraging NDI Signals:** When possible, opt for NDI (Network Device Interface) signals for video transport. NDI offers high efficiency and quality over IP networks, making it well-suited for scalable live production environments. This approach enables the easy addition of sources and destinations without significantly impacting the core production infrastructure.

Segmenting Network Traffic

- **VPC Configuration:** Use Amazon VPC to segment different types of network traffic, such as live feeds, monitoring data, and management operations. This segmentation facilitates more efficient traffic management and enhances security by isolating production elements from each other.
- **Traffic Control and Security:** Implement traffic flow controls and security measures within your VPC to ensure that critical live production traffic is prioritized and protected. Utilize network ACLs and security groups to enforce these policies.

Deployment for High Availability

- **Multi-AZ and Multi-Region Deployment:** As detailed in the "Reliability and Resilience" chapter, deploying your TriCaster Vector environment across multiple Availability Zones or even regions can significantly enhance the availability and fault tolerance of your live productions. This approach ensures that even in the event of a zone or region failure, your production can continue without significant disruption.

6.6 Conclusion

Optimizing your TriCaster Vector deployment on AWS for live cloud production is a dynamic process that requires careful consideration of computing, memory, storage, and network settings.

The recommendations provided in this guide offer a robust starting point designed to enhance performance efficiency and ensure the delivery of high-quality live broadcasts. However, it's crucial to recognize that the one-size-fits-all approach does not apply to the diverse and complex nature of live production workflows.




The unique requirements of your specific pipeline workflows might necessitate adjustments beyond these general recommendations. To ensure your deployment is not only configured for optimal performance but also aligned with your specific production needs, broadcasters are encouraged to undertake the following actions:

- **Perform Adequate Testing:** Comprehensive quality and load testing, tailored to your specific production workflows, are indispensable. Such testing procedures help in pinpointing any potential bottlenecks or performance issues that could impede your live production. Through rigorous testing, you can verify that the recommended settings provide the support your live production operations require, making adjustments as necessary to fine-tune your setup.
- **Monitor and Adjust:** The dynamic nature of live broadcasting means that your deployment's performance needs will evolve. Continuously monitor the performance of your setup, leveraging tools like Amazon CloudWatch to gain insights into how your resources are being utilized in real-time. Be prepared to make timely adjustments to computing, memory, storage, and network settings in response to changing production demands or as you scale your operations. This proactive approach ensures that your deployment remains efficient, resilient, and capable of delivering the seamless live broadcast experience your audience expects.

7 Reliability and Resilience

In the demanding world of live cloud production, reliability and resilience are not just objectives but necessities. The ability to deliver uninterrupted, high-quality broadcasts directly influences audience satisfaction and trust.

Central to leveraging AWS's capabilities effectively is understanding the Shared Responsibility Model, which outlines the division of resilience duties between AWS and the customer. While AWS ensures the resilience of the cloud infrastructure, customers are responsible for implementing resilience in the cloud—designing fault-tolerant applications, managing data backups, and planning for disaster recovery. This delineation emphasizes the customer's role in actively managing and configuring their resources to achieve the desired levels of reliability and resilience.



For live cloud productions using TriCaster Vector, this means meticulously planning resource usage, employing AWS services strategically for high availability, and continuously monitoring system performance. By embracing these responsibilities, broadcasters can ensure their live events run smoothly, maintaining the high standards expected by their audiences.

Designing for Reliability

Designing for reliability in live cloud production involves adhering to principles that ensure your architecture can sustain operations under varying conditions without compromising broadcast quality.



The foundation of a reliable architecture lies in embracing high availability design patterns, fault tolerance, and redundancy strategies.

- **High availability** is achieved by designing systems that are resilient to failures and capable of maintaining operational continuity, which is crucial for live event broadcasting. This involves deploying resources across multiple Availability Zones, ensuring that if one zone experiences issues, others can seamlessly take over.
- **Fault tolerance** requires systems to continue functioning even when components fail. For TriCaster Vector users, this means implementing redundant instances and leveraging services with Multi-AZ deployments for storage to ensure data integrity and availability.

Redundancy strategies further reinforce reliability by duplicating critical components, such as streaming servers and storage, to eliminate single points of failure. By incorporating these principles into your TriCaster Vector deployment, you create a robust environment that maximizes uptime and delivers a consistent, high-quality viewing experience to your audience.

Resilience Trade-offs

While designing for reliability it's equally important to recognize and navigate the trade-offs associated with enhancing resilience. These trade-offs — cost and effort, complexity, operational burden, and the balance between consistency and latency — play a critical role in decision-making processes and ultimately impact the architecture's effectiveness and efficiency.

- **Cost and Effort:** Implementing high availability and fault tolerance mechanisms, such as multi-AZ deployments and redundant systems, inevitably incurs higher costs and requires more upfront planning and effort. Customers must weigh the financial implications against the potential cost of downtime to determine the optimal level of investment in resilience.
- **Complexity:** As redundancy and fault tolerance features are added to an architecture, the complexity of the system increases. This complexity can introduce challenges in maintenance and configuration, requiring specialized skills and potentially leading to higher operational overhead.
- **Operational Burden:** The effort to manage and operate a highly resilient system, particularly one that spans multiple Availability Zones or regions, can significantly increase. Automated scaling, health checks, and failover processes must be meticulously managed to ensure seamless operation, adding to the operational workload.
- **Consistency and Latency:** Ensuring data consistency across geographically dispersed systems can introduce latency, particularly in distributed databases or when synchronizing stateful applications. Balancing the need for immediate data



availability with acceptable latency levels is crucial, especially in live production environments where real-time performance is key.

Incorporating Recovery Objectives into Resilience Design

An essential complement to understanding the trade-offs in resilience design for live cloud production is defining clear recovery objectives. Establishing your Recovery Point Objective (RPO) and Recovery Time Objective (RTO) provides a concrete foundation upon which to build and measure the resilience of your architecture.


- **Recovery Point Objective (RPO):** The RPO represents the maximum age of data that must be recovered from backup storage for normal operations to resume after a fault or failure. In the context of live cloud production, a low RPO is crucial for minimizing data loss, ensuring that any recorded content, configurations, or state information lost during an incident is limited to a tolerable window.
- **Recovery Time Objective (RTO):** The RTO is the targeted duration of time within which a business process must be restored after a disaster or disruption to avoid unacceptable consequences. For live broadcasting operations using TriCaster Vector, the RTO dictates how quickly your production can recover and resume, which is critical for maintaining audience engagement and meeting contractual obligations.

Setting these recovery objectives at the outset guides the design of your system's resilience. A desired low RPO and RTO might necessitate more sophisticated and potentially costlier redundancy and failover mechanisms, such as real-time data replication across multiple Availability Zones. Conversely, understanding the acceptable thresholds for data loss and downtime can help in making informed decisions that balance resilience with cost, complexity, and operational overhead.

7.1 Disaster Recovery Strategies for TriCaster Vector

Disaster Recovery (DR) is a critical component of any live cloud production pipeline, including deployments using TriCaster Vector on AWS. AWS categorizes DR strategies into four main approaches, each offering different levels of cost, complexity, and readiness. For TriCaster Vector users, selecting the right DR strategy involves considering the entire live cloud production pipeline, from network contribution through to distribution.

Here, we outline the DR options suitable for TriCaster Vector, emphasizing the need for comprehensive planning and regular testing.



In this section we delve into the "Backup and Restore" and "Warm Standby" methodologies, providing a framework that not only secures your operations against potential disruptions but also facilitates rapid recovery.

While the "Pilot Light" approach is a feasible strategy, it can be derived and customized by users based on the principles and practices outlined in the aforementioned strategies, taking into account the intricacies of the overall pipeline.

Currently, the "Multi-Site Active/Active" strategy is in the development phase, with future insights to be shared through the [AWS Virtual Live Remote Production \(vLRP\) Partner Acceleration Initiative](#). This initiative aims to innovate and establish comprehensive solutions tailored for media environments, signifying our commitment to advancing disaster recovery capabilities.

Strategy	Overview	Application to TriCaster Vector
Backup and Restore	Fundamental DR strategy with periodic system, configuration, and data backups. Restoration reinstates the environment.	Implement routine backups of TriCaster Vector settings and AWS-stored data (EBS, S3). Suitable for scenarios with acceptable downtime recovery limits.
Pilot Light	Keeps critical system elements operational in the cloud, ready to scale up during a disaster. Offers readiness at a reduced cost.	Maintain an always-updated, minimal version of TriCaster Vector on AWS, including application configurations and replicated media. Allows for faster recovery than backups alone.
Warm Standby	A scaled-down, operational version of the full environment is maintained, allowing for quick scaling and reduced RTO.	Operate a secondary TriCaster Vector setup in standby within another AWS Region or Availability Zone, ready to take over swiftly in case of disaster.



Multi-site Active/Active	Advanced strategy with the production pipeline operating simultaneously across multiple locations for uninterrupted service.	Under development: Aiming to support parallel TriCaster Vector deployments across various AWS Regions for seamless failover and resilience. (Note: Currently not fully supported.)
---------------------------------	--	--

7.2 Backup and Restore

The Backup and Restore approach is pivotal for mitigating data loss or corruption and ensuring resilience against regional disasters in live cloud productions using TriCaster Vector. By replicating data across AWS Regions or safeguarding against resource scarcity, such as GPU availability, this strategy underpins a robust disaster recovery plan.

Key Components:

1. Critical Data and Configuration Elements:

- **EBS Volumes for TriCaster Vector:** Store OS, software configurations, and operational metadata, including user settings, recordings, raw footage, and digital assets used in live productions should be regularly backed up.
- **License Server Volumes:** Ensure all licensing information is backed up to facilitate swift recovery.
 - **The backup strategy relies on EBS volume snapshots:** Owing to the Wibu security feature, the snapshot must be restored to the original instance where it was created, maintaining the same instance ID. This procedure is designed to ensure recovery from software crashes exclusively.
- **License Server Rehosting Licenses:** Leverage rehosting capabilities to transfer licenses from the License Server back to the License Portal, allowing to transfer existing Licenses to a new installation.
- **Amazon S3:** Use for optional storage of deployment scripts or media files, enhancing data durability.
- In addition to user data, be sure to also back up code and configuration, including Amazon Machine Images (AMIs) you use to create Amazon EC2 instances.

2. Automating Backups:


- Leverage AWS services such as Amazon EBS snapshots and Amazon S3 for automated, scheduled backups, incorporating incremental backups to optimize storage.



- How often you run your backup will determine your achievable recovery point (which should align to meet your RPO).

3. Secure Storage and Management:

- Utilize AWS Backup for centralized backup management of EBS volumes and enable Amazon S3 Cross-Region Replication (CRR) for continuous asynchronous copying. Implement S3 object versioning as an additional layer of protection against human error.
- **Lifecycle Management:** Apply lifecycle policies to manage backup retention, automatically archiving or deleting old backups based on predefined rules to optimize storage costs.




4. Restoration Process and Testing:

- **Infrastructure and Pipeline Restoration:** Ensure the entire infrastructure, including TriCaster Vector and License Server, can be redeployed in a new Availability Zone or Region following a disaster.
- To enable infrastructure to be redeployed quickly without errors, you should always deploy using infrastructure as code (IaC) using services such as Viz Now, AWS CloudFormation or the AWS Cloud Development Kit (AWS CDK). Without IaC, it may be complex to restore workloads in the new recovery



Zone or Region, which will lead to increased recovery times and possibly exceed your RTO.

- Develop clear, detailed restoration procedures and regularly test the restoration process to validate RPO and RTO targets.



5. Documentation and Training:

- Maintain comprehensive documentation of the backup and restore process and train all team members on the strategy to ensure readiness.

6. Monitoring and Alerts:


- Utilize AWS CloudWatch to monitor backup health and configure alerts for any issues that arise, ensuring prompt attention to failed backups or other critical concerns.

Alternative Considerations:

- **Amazon Machine Images (AMIs):** Consider backing up EC2 instances as AMIs for a more holistic recovery solution, allowing for instance restoration across regions. AWS Backup enhances EC2 backup capabilities by including metadata essential for comprehensive restoration.

7.3 Warm Standby

The Warm Standby disaster recovery (DR) strategy for TriCaster Vector involves maintaining an exact replica of the primary region's assets in a DR region, ready to handle production immediately without the need to activate servers. This approach ensures high availability and rapid recovery in the event of a disaster.



System Configuration Alignment

- **Pre-Production Alignment:** Before initiating production, it's crucial to synchronize the system configurations of TriCaster Vector instances across both the primary and DR regions. This ensures consistency and seamless operation across environments.
- **Session Backup and Restore:** Utilize TriCaster Vector's session backup and restore functionality to streamline the alignment process. This feature facilitates easy replication of settings, ensuring both systems mirror each other's configurations accurately.

Multi-AZ File System Setup

- **File System for Clip Playback:** To efficiently manage and share resources such as clips between multiple Vector instances, setting up a Multi-AZ file system is recommended.
- **Amazon FSx for Windows File Server:** Opt for Amazon FSx with a Multi-AZ file system deployment for its robustness and compatibility with Windows-based applications. This setup enhances resource sharing and accessibility across regions.



- **Throughput Testing:** Conduct thorough testing to determine the required throughput for recording and playback, tailoring the file system's performance to meet your specific workflow demands.

Operational Considerations

- **NICE DCV for Active System Operations:** Regular operations are conducted via a NICE DCV connection to the active system, allowing for real-time production management and monitoring.
- **Clip Replication:** Created clips are automatically replicated to the passive zone via Amazon FSx for Windows File Server, ensuring that all resources are up-to-date and available for immediate use in the DR region.
- **Rapid Switching:** Operators can maintain a second NICE DCV connection to the standby system, enabling a swift transition to the DR region without significant downtime.
- **Content Distribution Network Reconfiguration:** The Business Continuity Planning (BCP) procedure should account for necessary adjustments to the video origin settings for the content distribution network (CDN) during the switch to the standby zone. This ensures that the distribution of content remains uninterrupted and reaches the intended audience seamlessly.

7.4 Multi-Site Active/Active

The Multi-Site Active/Active strategy is not yet supported for TriCaster Vector however, significant strides are being made towards its realization through the AWS Virtual Live Remote Production (vLRP) Partner Acceleration Initiative. This collaborative effort, involving a consortium of leading vendors, is dedicated to pioneering innovative solutions that cater specifically to the demands of media environments.

The solution focuses on developing a robust architectural framework that allows for real-time synchronization between dual TriCaster Vector instances. This synchronization ensures seamless mirroring of commands from a primary (master) node to a secondary (slave) node, facilitating a fully redundant backup video output identical to the primary feed.

The envisioned architectural solution aims to:

- **Maintain Synchronization:** Keep two TriCaster Vector instances in perfect sync, ensuring that all commands and operations performed on the master node are instantly replicated on the slave node.



- **Ensure Redundancy:** Provide a fail-safe mechanism, enabling the slave node to take over immediately and seamlessly should the primary node encounter any issues, with no disruption to the live video output.
- **Support Seamless Operations:** Enable production teams to operate in a truly active/active environment, where both nodes are fully operational offering unparalleled resilience and uptime.

As the AWS vLRP Partner Acceleration Initiative progresses, we anticipate releasing detailed guidelines and technical specifications for implementing the Multi-Site Active/Active strategy with TriCaster Vector.

7.5 Monitoring and Incident Response

Ensuring the reliability and resilience of TriCaster Vector deployments in live cloud production environments requires diligent monitoring and a proactive incident response strategy. This section outlines approaches to maintain operational integrity and quickly address any issues that arise.

Monitoring

Effective monitoring acts as the foundation for understanding and maintaining the health of your TriCaster Vector deployment. It enables you to detect issues before they impact your live production, ensuring continuous availability and optimal performance.

- **Refer to the "Operational Excellence" Chapter:** For comprehensive insights into establishing a robust monitoring framework, please refer to the "Operational Excellence" chapter. It covers best practices for leveraging AWS services and tools to monitor your deployment continuously.
- **Performance Efficiency Insights:** Additionally, the "Performance Efficiency" chapter provides valuable information on monitoring the performance of your AWS resources. Understanding and applying these principles ensures that your deployment is not only reliable but also operates at peak efficiency.

Incident Response

When incidents occur, having a well-defined incident response plan is crucial for minimizing impact and restoring service quickly. For TriCaster Vector and live cloud production, consider the following recommendations:

1. **Pre-define Response Protocols:** Establish clear protocols for different types of incidents, including performance degradation, network issues, or security breaches. Each protocol should outline steps for identification, assessment, containment, and resolution.



2. **Automate Responses Where Possible:** Utilize AWS services like AWS Lambda in conjunction with Amazon CloudWatch alerts to automate certain response actions. For example, automatically restarting an instance or scaling resources in response to specific triggers can reduce downtime.
3. **Regularly Update and Test Incident Response Plans:** As your live production environment evolves, so should your incident response plans. Regular testing through simulations or drills will help ensure your team is prepared and response mechanisms are effective.
4. **Communicate Effectively During Incidents:** Maintain transparent communication with stakeholders during an incident. Provide regular updates on the status, impact, and expected resolution time.
5. **Post-Incident Analysis:** After resolving an incident, conduct a thorough analysis to identify the root cause and lessons learned. Use this information to refine your monitoring and response strategies, preventing future occurrences.
6. **Leverage Vizrt Global Support:** For incidents related to TriCaster Vector, don't hesitate to reach out to Vizrt Global Support for expert assistance. Their knowledge can be invaluable in quickly addressing product-specific issues.

7.6 Conclusion

Implementing a DR strategy for TriCaster Vector requires a holistic view of your live cloud production pipeline. Each phase of the broadcast workflow, from network contribution to distribution, must have appropriate DR strategies in place. Regular assessment and testing of your DR plan are essential to ensure its effectiveness. AWS Resilience Hub can be an invaluable tool in this process, offering continuous validation of your AWS workloads' resilience and ensuring you are likely to meet your RTO and RPO targets.

8 Operational Excellence

Achieving operational excellence is crucial for broadcasters utilizing TriCaster Vector within AWS environments, enabling seamless, efficient delivery of high-quality live content. This concept extends beyond technical prowess, encompassing strategic planning, execution, and continuous improvement to optimize performance, reliability, and scalability. It necessitates a comprehensive understanding of TriCaster Vector and AWS services, focusing on integrating these platforms to develop resilient and flexible production workflows.

Operational excellence involves leveraging best practices in cloud architecture, employing cost-effective deployment strategies, maintaining robust security measures, and adopting proactive monitoring and management. It signifies a commitment to continuous evaluation and adaptation, ensuring that live broadcasting operations remain efficient, secure, and capable of meeting evolving demands.

This introduction highlights the essence of operational excellence for TriCaster Vector deployments, emphasizing the importance of efficiency, reliability, and innovation in live cloud production. It sets the groundwork for exploring strategies and practices that foster an operationally excellent broadcasting environment.

8.1 Designing for Operational Efficiency

Designing for operational efficiency involves architecting your cloud environment in a way that maximizes performance and minimizes waste, thereby enhancing the overall quality of live broadcasts. Here's how to approach designing for operational efficiency with TriCaster Vector deployments:

Leveraging AWS Cloud Architecture Best Practices

- **Well-Architected Framework:** Align your TriCaster Vector deployment with the AWS Well-Architected Framework, focusing on the operational excellence pillar. This involves understanding and implementing AWS best practices for security, reliability, performance efficiency, cost optimization, and sustainability.

Enhancing Performance and Scalability

- **Content Delivery Network (CDN):** Implement Amazon CloudFront to distribute your live broadcast content globally. A CDN reduces latency and improves the viewing experience for your audience by caching content at edge locations closer to users.

Cost Optimization



- **Right-Sizing Resources:** Regularly review and adjust the size of your AWS resources to match the demand. Use AWS Trusted Advisor and Cost Explorer to identify underutilized resources and potential savings opportunities.
- **Reserved Instances and Savings Plans:** For predictable workloads, consider purchasing Reserved Instances or committing to AWS Savings Plans to reduce costs significantly compared to On-Demand pricing.

High Availability and Fault Tolerance

- **Multi-AZ Deployments:** Deploy TriCaster Vector environment across multiple Availability Zones to ensure high availability and fault tolerance. This approach protects your live production against data center failures.
- **Backup and Disaster Recovery:** Establish robust backup and disaster recovery plans using AWS technologies. Regular testing of these plans ensures that you can quickly recover from unforeseen incidents.

Security and Compliance

- **Data Protection:** Implement encryption, both at rest and in transit, using AWS services like Amazon EBS, Amazon S3 and AWS Key Management Service (KMS) to protect your content and comply with industry regulations.
- **Access Management:** Use AWS Identity and Access Management (IAM) to enforce the principle of least privilege, ensuring that only authorized users have access to your AWS resources.

8.2 Deployment Strategies for TriCaster Vector

Deploying TriCaster Vector in the AWS cloud requires a strategic approach that balances efficiency, scalability, and reliability. By adopting specific deployment strategies, broadcasters can ensure their live production environments are optimized for performance and cost-effectiveness. Here are key deployment strategies for TriCaster Vector:

Infrastructure as Code (IaC)

- **Overview:** IaC is a key practice that automates the provisioning and management of your cloud infrastructure using code, rather than manual processes. It ensures consistency, reduces the potential for human error, and speeds up deployment and scaling processes.
- **Implementation with TriCaster Vector:** Utilize AWS CloudFormation or the AWS Cloud Development Kit (AWS CDK) to define your TriCaster Vector environment as code. This includes networking configurations, EC2 instances, storage, and any other AWS services required for your production. IaC enables rapid deployment across



different environments or regions and simplifies the process of replicating your setup for disaster recovery or scaling purposes.

Utilizing Viz Now for Streamlined Setup

- **Overview:** Viz Now is a deployment tool offered by Vizrt that simplifies the process of setting up TriCaster Vector within AWS. It provides templates and workflows designed to get your live production environment up and running quickly.
- **Benefits:** By leveraging Viz Now, broadcasters can significantly reduce the setup time and complexity associated with deploying TriCaster Vector. It ensures that best practices are followed, and the environment is configured for optimal performance.

8.3 Monitoring and Management

Effective monitoring and management are pivotal for maintaining operational excellence in TriCaster Vector deployments on AWS. These practices provide insights into system performance, resource utilization, and potential issues, allowing for proactive adjustments and ensuring high-quality live productions. Here's how to approach monitoring and management for TriCaster Vector:

Implementing Comprehensive Monitoring

- **AWS CloudWatch:** Utilize Amazon CloudWatch for comprehensive monitoring of your TriCaster Vector environment. CloudWatch collects and tracks metrics, collects and monitors log files, and sets alarms. Use it to monitor AWS resources such as Amazon EC2 instances and Amazon EBS volumes, ensuring they are performing as expected. Set up alarms for key performance indicators like CPU utilization, network bandwidth, and error rates to receive immediate notifications of potential issues.
- **Real-Time Analytics:** Implement CloudWatch Dashboards to create customizable home pages for monitoring the health and performance of your TriCaster Vector environment in real-time. This enables quick access to metrics that matter most, aiding in swift decision-making.

Proactive Management Practices

- **AWS Sizing:** Ensure your TriCaster Vector EC2 deployment instance capacity is right-sized to maintain steady, predictable performance at the lowest possible cost.
- **Resource Optimization:** Regularly review CloudWatch metrics and logs to identify underused resources that can be downsized, stopped, or terminated to save costs without impacting production quality. Similarly, identify bottlenecks where scaling up resources could improve performance.

Setting Up Alerts and Notifications




- **Proactive Alerts:** Configure CloudWatch Alarms to send alerts via Amazon SNS when specific thresholds are breached. This could include alerts for high CPU utilization, low disk space, or network bottlenecks, enabling rapid response to potential issues before they impact your live broadcast.
- **AWS Budgets Alerts:** Use AWS Budgets to set custom cost budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. This is crucial for managing costs effectively and avoiding unexpected expenses.

Security Monitoring and Compliance

- **AWS CloudTrail:** Leverage AWS CloudTrail to log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. CloudTrail provides visibility into user and resource activity, making it easier to ensure compliance with internal policies and regulatory standards.
- **IAM Access Analyzer:** Use IAM Access Analyzer to analyze resource permissions and identify policies that allow access to your resources from outside your AWS account, ensuring that only intended users have necessary access.
- **VPC Flow Logs:** It's recommended to enable VPC Flow Logs to capture information about the IP traffic going to and from network interfaces in your Virtual Private Cloud (VPC). This feature allows you to capture and inspect network traffic patterns, identify anomalies, and ensure that network activity aligns with your security and operational policies. VPC Flow Logs serve as a crucial tool for diagnosing traffic behaviors, investigating security concerns, and achieving a comprehensive understanding of your network traffic for compliance and auditing purposes.

8.3.1 Pipeline Monitoring Hub

To enhance live service management within TriCaster Vectar deployments, integrating a dedicated EC2 instance into the backing services infrastructure is recommended. This specialized instance, known as the Pipeline Monitoring Hub, serves as the operational core for the Live Service Manager, offering comprehensive health checks and monitoring capabilities to ensure the seamless operation of the entire live production pipeline.



Recommended EC2 Instance Specification

- **Instance Type:** For tasks requiring real-time video monitoring and remote command execution within the live production workflow, select an instance with robust compute capacity and network performance. The **g4dn.2xlarge** instance type is recommended as a starting point, though the final choice should be tailored to the specific tools and requirements of your production pipeline.

Tools and Functionalities

- **Real-Time Video Monitoring:** For NDI-based pipelines, ensure the installation of NDI Tools and the NDI Analysis Tool on your system to facilitate real-time video monitoring. Integrate Prometheus with NDI and SRT Node Exporter to enable comprehensive monitoring through Grafana, further enhanced by integration with AWS CloudWatch.
- **Remote Command Execution:** Enhance the Pipeline Monitoring Hub by integrating essential software that enables the Live Service Manager to perform remote commands throughout the live production workflow. This includes tools akin to Windows IT Admin Center, which supports the execution of remote PowerShell commands via SSH, and AWS System Manager, offering capabilities such as remote CLI, fleet-wide remote access directly from the AWS console, as well as remote installation via the AWS console or CLI. This enhancement streamlines the process of



making quick operational changes while minimizing interference with other operators' activities.

IAM Role and Access Management

- **Specific IAM Role:** Create a unique IAM role for the EC2 instance, granting it permissions to monitor and manage the live production infrastructure effectively. This role should include access to AWS CloudWatch for monitoring and any other AWS services integrated into your live production environment monitoring.
- **Access Control:** Enhance security measures by implementing rigorous access controls, ensuring that only authorized personnel, such as the Live Service Manager, have access to the Pipeline Monitoring Hub. Utilize AWS IAM policies to meticulously define and enforce these access rules. Additionally, integrate AWS GuardDuty for continuous monitoring and detection of threats, further strengthening the security and integrity of your operations.

Implementation Strategy

- **Provisioning:** Utilize AWS CloudFormation or the AWS Cloud Development Kit (AWS CDK) for provisioning this EC2 instance as part of your backing services. Infrastructure as Code (IaC) ensures consistent and efficient deployment.
- **Configuration and Setup:** Post-provisioning, configure the instance with essential tools and software for effective live service management. Standardize this setup to guarantee replicability across different environments or regions as needed. Facilitate access to various VPCs via VPC peering or other means for seamless P2P connectivity.
- **Monitoring and Maintenance:** Keep the installed tools and functionalities up to date with the latest features and security patches. Regularly review the instance's performance, making necessary adjustments based on the evolving demands of your live production.

8.4 Planning EC2 GPU-based Instance Provisioning

In the pursuit of operational excellence, it's essential to address the provisioning of EC2 GPU-based instances with meticulous planning. GPU instances are critical for in a live cloud productions and can be subject to availability constraints due to high demand.

Understanding GPU Instance Availability

- **Availability Constraints:** AWS offers EC2 GPU-based instances that are ideal for compute-intensive tasks required in live broadcasting. However, due to their specialized nature and high demand, these instances can sometimes be scarce, with no guaranteed availability on-demand across all regions.



- **Impact on Live Production:** For TriCaster Vector deployments, the reliance on GPU instances for real-time video processing means that any scarcity can directly impact the ability to deliver live content. Planning ahead is crucial to avoid disruptions.

Strategies for Ensuring GPU Instance Availability

- **Capacity Reservations:** AWS provides the option to reserve capacity for EC2 instances ahead of time. By pre-provisioning the required GPU-based instances, you secure the necessary resources, ensuring they are available when your live production goes on air. This approach is highly recommended for broadcasters with predictable production schedules and can mitigate the risk of instance scarcity.
- **Utilizing Savings Plans or Reserved Instances:** While different from capacity reservations in terms of physical resource allocation, committing to Savings Plans or purchasing Reserved Instances can offer both cost savings and a higher priority for instance allocation compared to on-demand users. This financial commitment can indirectly support your case for capacity reservation with AWS.
- **Diversifying Instance Types and Regions:** To mitigate the risk of scarcity, consider designing your live production environment to be flexible across different GPU instance types and AWS regions. This diversification strategy increases the chances of securing the necessary compute resources by tapping into broader availability pools.

Pre-Provisioning Capacity

- **Engaging with AWS Support:** Early engagement with AWS support can facilitate the process of pre-provisioning GPU-based instances. AWS account managers can provide insights into availability trends and assist in securing the necessary capacity for your specific requirements.
- **Planning for Peak Events:** For live events expected to draw significant audiences, or during peak industry periods, planning and reserving capacity well in advance becomes even more critical. Assess historical data and future event schedules to determine when increased demand for GPU instances is likely and pre-provision accordingly.

8.5 Designing the Contribution Network

Designing a robust contribution network is essential to ensure uninterrupted live video and audio production. The network must have sufficient bandwidth to handle the upload and download of live feeds, multi-viewer outputs for monitoring, and remote display protocols without degradation in quality.

Step 1: Calculate Upload and Download Requirements



- **Live Feeds:** Start by calculating the bandwidth needs for your live feeds. Account for both video and audio bitrates, as high-definition content and high-quality audio significantly increase bandwidth requirements. Sum up the bitrates for all live feeds to determine total upload and download needs.

Step 2: Consider Multi-Viewer Output for Monitoring

- **Monitoring Screens:** Add the bandwidth requirements for multi-viewer outputs used for production monitoring. Higher resolution and frame rates increase bandwidth demands. Aggregate the bandwidth for all screens to calculate the total requirements for multi-viewer monitoring.

Step 3: Consider Nice DCV Requirements

- **Remote Display Protocol:** Nice DCV, used for high-performance remote displays, adds to the bandwidth requirements. Assess the number of screens and their resolutions and frame rates to estimate the bandwidth needed for Nice DCV.

Step 4: Consider Transport Stream Requirements

- **Codec and Error Correction:** The choice of video codec and the application of error correction techniques influence bandwidth needs. Opt for codecs like H.264, which balance video quality with bandwidth efficiency. Factor in additional bandwidth for error correction to maintain stream quality over less reliable networks.

Step 5: Plan for Contingency

- **Additional Bandwidth:** Incorporate a contingency buffer of 20-30% on top of your estimated bandwidth requirements to accommodate unexpected demand spikes or network fluctuations, ensuring a reliable live stream.

Ensuring EC2 Networking Capability

- **Instance Networking Capability:** Verify that the chosen EC2 instances support the required networking performance to manage the data flow. AWS offers instances with enhanced networking capabilities, suitable for high-throughput live video applications. Ensure your instance selection aligns with the network performance needed for your live feeds, multi-viewer output, and Nice DCV requirements.

This example demonstrates how to estimate bandwidth requirements for a live production setup using NDI (Network Device Interface) with 1920x1080 60P NDI|HX H.264 codec. The scenario involves 8 NDI video sources and an operator monitoring the production through a single screen and a multi-viewer feed.

Component	Profile	Mbps	Number of Streams	Total Mbps
<i>Upload - to the cloud</i>				



Video Audio Source	1920*1080 60P NDIIHX H.264	11	8	88
+20% contingency				106
<i>Download - from the cloud</i>				
NICE DCV	1920*1080 60P NDIIHX H.264	11	2	22
Multi-Viewer	1920*1080 60P NDIIHX H.264	11	1	11
+20% contingency				40

8.6 Continuous Improvement Practices

Continuous improvement is an essential principle for achieving operational excellence and it involves an ongoing cycle of evaluating, optimizing, and enhancing live production workflows to meet evolving requirements and leverage new technological advancements.

Embracing a Culture of Feedback and Iteration

- **Feedback Loops:** Establish mechanisms to gather feedback from all stakeholders involved in the live production process, including technical staff, content creators, and the audience. Use this feedback to identify areas for improvement in both the technical infrastructure and the content delivery experience.
- **Iterative Development:** Encourage regular reviews of operational practices and the adoption of an iterative approach to implementing changes, allowing for incremental improvements over time.

Leveraging AWS Tools for Ongoing Evaluation

- **AWS Well-Architected Tool:** Utilize the AWS Well-Architected Tool to regularly assess your TriCaster Vector deployment against AWS best practices. This tool provides insights into potential issues and offers guidance on optimizing your architecture for performance, security, and cost.
- **AWS Trusted Advisor:** Engage with AWS Trusted Advisor for real-time recommendations on improving your AWS environment. Trusted Advisor checks can help identify opportunities for reducing costs, enhancing security, and improving fault tolerance.

9 Cost Optimization

In the evolving landscape of cloud deployments, cost optimization represents a pivotal aspect of managing and scaling digital broadcasting environments effectively. As organizations transition to cloud-based solutions like TriCaster Vector for live event broadcasting, the ability to optimize costs without compromising on performance or reliability becomes increasingly important.

The deployment of TriCaster Vector within AWS harnesses the power of cloud computing to elevate live production workflows, offering broadcasters a versatile and robust platform for their operations. However, the true potential of this combination is realized when cost-saving measures are actively pursued and implemented. By leveraging AWS's comprehensive suite of services in tandem with TriCaster Vector, organizations can benefit from significant cost efficiencies through several mechanisms:

- **Scalability:** Dynamically adjusting resource allocation based on demand ensures that you pay only for the resources you use, preventing overprovisioning and underutilization.
- **Reserved Instances and Savings Plans:** For predictable workloads, committing to reserved instances or adopting savings plans can lead to substantial reductions in operational costs over time.
- **Right-Sizing:** Regularly analyzing and adjusting the size of your instances to match the workload requirements can lead to immediate cost savings.
- **Data Transfer Optimization:** Minimizing data transfer costs by keeping data transfer within the AWS network as much as possible and using content delivery networks efficiently.

This chapter aims to guide users through the intricacies of cost optimization when deploying TriCaster Vector on AWS, highlighting strategies, tools, and best practices designed to maximize cost efficiency while ensuring a high-quality broadcasting experience.

9.1 Budgeting and Cost Management Strategies

As live broadcasting evolves into more cloud-centric workflows, understanding and anticipating the various cost implications becomes crucial. This section outlines key considerations and strategies to ensure financial efficiency while maintaining high production standards.

Considerations for Cost Management

- **Broadcast Schedule and Resource Utilization:** Align your cloud resource usage with your broadcast schedule. Consider not only the live event times but also the required setup, rehearsal, and post-production times. Efficient scheduling ensures you're not paying for resources outside of active production hours.
- **Signal Contribution and Distribution Requirements:** The quality of the broadcast signal and the codecs used directly impact bandwidth and storage costs. Higher quality signals offer better viewer experiences but at higher data transfer and storage costs. Similarly, managing multiple video origins increases complexity and potentially costs. It's essential to balance these factors based on your production goals and budget constraints.
- **Operator Location and Preview Signals:** The geographical location of your operators can influence data transfer costs, especially if preview signals or control surfaces are delivered over long distances. Optimize preview signal quality to ensure operational efficiency without unnecessarily inflating costs.
- **Content Storage and Archiving:** Post-production content management, including storage and archiving, should be planned in advance. Consider the hours of production and the quality of the material to estimate storage needs accurately. Implementing data lifecycle policies can help transition older content to more cost-effective storage solutions over time.

Strategies for Budgeting and Cost Management

1. **Detailed Planning and Forecasting:** Start with a comprehensive plan that considers all aspects of your live production, from pre-production activities to content archiving. Use historical data and benchmarks to forecast resource needs and associated costs accurately.
2. **Leverage AWS Pricing Models:** Choose the most suitable AWS pricing models (On-Demand, Reserved Instances, Savings Plans) based on your production schedule and resource utilization patterns. Mixing and matching different pricing models can optimize costs without compromising flexibility.
3. **Optimize Data Transfer and Storage:** Use services like Amazon CloudFront to reduce data transfer costs and improve audience experiences. For storage, consider the trade-offs between immediate access (Amazon EBS, Amazon EFS) and long-term archiving costs (Amazon S3, Amazon S3 Glacier).
4. **Continuous Monitoring and Optimization:** Employ AWS tools like Cost Explorer and Budgets to monitor spending and usage patterns in real-time. Regular reviews allow for adjustments to be made proactively, avoiding budget overruns.



5. **Collaborate with AWS Support and Use Cost Management Tools:** Engage with AWS support for insights into cost-saving opportunities and best practices. Utilize AWS cost management tools for granular tracking and optimization of your spending.

9.2 Cost-Efficient Resource Utilization

Efficient resource utilization is a cornerstone of optimizing cloud deployments for live event broadcasting. It entails scaling resources to align precisely with production demands, thereby avoiding the common pitfall of over-provisioning, which can lead to unnecessary costs without yielding performance benefits.

Scaling Resources to Match Production Demands

- **Right-Sizing Instances:** Start with the recommended instance size, as indicated in the Performance Efficiency chapter, which suggests beginning at a g4dn.2xlarge and scaling up to a g5.8xlarge based on needs. It's crucial, however, to tailor your selection based on specific production demands.
- **Quality Assurance Testing:** Before finalizing your instance selection, conduct comprehensive Quality Assurance Testing, including Performance and Stress Testing as outlined in the "Quality Assurance and Testing" chapter. This testing will validate whether the chosen instance types and sizes meet the demands of your live broadcasting workload without under or over-provisioning.
- **Cost Optimization Across Regions:** Be aware of the cost differences between AWS regions. Deploying resources in regions with lower costs can lead to significant savings, provided that it doesn't negatively impact latency or compliance with data residency requirements.

Eliminating Inefficiencies

- **Identify Underutilized Resources:** Regularly review your resource utilization metrics to identify and decommission underutilized instances or downgrade them to more cost-effective sizes.

9.3 Understanding AWS Pricing Models

AWS offers a variety of pricing models designed to provide flexibility and cost savings for different use cases. Understanding these models and their respective trade-offs is crucial for optimizing your cloud infrastructure costs while meeting your broadcasting requirements. Here's a concise overview of the main AWS pricing models:

On-Demand Instances



- **Description:** On-demand instances allow you to pay for computing capacity by the hour or second (minimum of 60 seconds) with no long-term commitments. This model offers maximum flexibility, allowing you to scale your resources up or down based on your needs without upfront payment.
- **Best for:** TriCaster Vector deployments with variable workloads, short-term, sporadic, or unpredictable requirements where the cost premium of On-Demand is justified by the flexibility it provides.
- **Trade-offs:** Offers the highest cost per compute capacity compared to other pricing models. Ideal for situations where the ease of use and flexibility outweigh the need for cost optimization.

Reserved Instances

- **Description:** Reserved Instances provide a significant discount compared to On-Demand pricing in exchange for committing to a specific instance type and usage (1-year or 3-year term) in a particular region.
- **Best for:** Deployments with stable and predictable workloads, where long-term planning is possible. Suitable for continuous TriCaster Vector operations that run throughout the year.
- **Trade-offs:** While offering considerable cost savings, Reserved Instances require upfront payment and long-term commitment, reducing flexibility to change instance types or regions.

Savings Plans

- **Description:** Savings Plans offer a more flexible model than Reserved Instances, providing savings in exchange for committing to a consistent amount of usage (measured in \$/hour) over a 1 or 3-year term. AWS offers two types of Savings Plans: Compute Savings Plans, which apply to any EC2 instance regardless of region, instance family, OS, or tenancy; and EC2 Instance Savings Plans, which provide the deepest discounts for specific instance families in a region.
- **Best for:** Users seeking cost savings with some level of flexibility. Compute Savings Plans are particularly suited for TriCaster Vector deployments where usage patterns might vary but overall compute spend remains consistent.
- **Trade-offs:** Savings Plans offer less flexibility than On-Demand but more than Reserved Instances. Users commit to a specific amount of spend, not instance configurations, allowing for some adjustments as needs change.

Leveraging AWS Pricing Models for TriCaster Vector Deployments



Choosing the right pricing model depends on your specific use case, budget, and operational flexibility. For TriCaster Vector deployments:

- Evaluate your usage patterns to determine if they're consistent enough for Reserved Instances or Savings Plans.
- Consider using On-Demand Instances for unpredictable spikes in demand or for short-term projects.
- Mix and match pricing models based on different workloads and environments within your deployment to optimize cost without sacrificing the necessary flexibility.

9.4 Storage Optimization

Effective management of storage resources is a critical component of cost optimization in cloud deployments, especially for applications like TriCaster Vector that handle extensive video content. AWS provides a range of storage solutions, each with unique features designed to help manage costs while meeting your storage needs. This section offers guidance on utilizing Amazon S3 and Amazon Elastic Block Store (EBS) efficiently, focusing on strategies to optimize storage costs without compromising accessibility or performance.

Amazon Elastic Block Store (EBS) Sizing and Snapshots

EBS provides persistent block storage volumes for use with EC2 instances. For TriCaster Vector deployments:

- **Right-Size EBS Volumes:** Carefully size your EBS volumes based on the expected hours of the recording material, considering both contingency needs and the quality of the material. The codec used for managing streams significantly impacts storage requirements.
- **Optimize EBS Throughput:** In addition to sizing your EBS volumes based on storage capacity needs, it's essential to consider the throughput requirements, especially given the number of feeds being recorded simultaneously. The Performance Efficiency chapter details how to evaluate and adjust EBS throughput to match the performance demands of your live broadcasts. This optimization ensures that your storage solution can handle the data transfer rates required by multiple, concurrent video streams without bottlenecks, thereby maintaining the integrity and quality of your recordings.
- **Leverage EBS Snapshots for Archiving:** For material that does not require real-time access and needs to be archived, consider moving it to EBS snapshots and then to S3. EBS snapshots are incremental backups that save money by storing only the changed blocks since the last snapshot.



Amazon S3 Efficiency and Lifecycle Policies

Amazon S3 is an excellent solution for storing and retrieving any amount of data, at any time, from anywhere on the web. For TriCaster Vector users, it offers a cost-effective way to store recorded material, raw footage, and produced content. To optimize costs:

- **Utilize S3 Lifecycle Policies:** Automate the transition of data to more cost-effective storage classes or archive data that is infrequently accessed but must be retained. For example, move older content to S3 Glacier or S3 Glacier Deep Archive for long-term storage at lower costs.
- **Delete Unused Data:** Regularly review and purge unnecessary data. S3 Lifecycle policies can also automate this process by deleting data that has reached the end of its lifecycle.

Best Practices for Storage Cost Optimization

- **Regularly Review Storage Needs:** Continuously monitor your storage utilization and adjust your strategy as requirements change.
- **Use AWS Cost Management Tools:** Tools like AWS Cost Explorer and the AWS Budgets can help track storage costs and usage, enabling informed decisions about where optimizations can be made.
- **Educate Your Team:** Ensure that your team understands the costs associated with storage options and the importance of following best practices for data management.

9.5 Optimizing Data Transfer Costs

Data transfer costs can significantly impact the overall budget of cloud-based live event broadcasting. By strategically managing data transfer, customers can minimize costs while maintaining or even enhancing the viewer experience.

Understanding AWS Data Transfer Costs

AWS charges for data transfer based on the amount of data moving out of AWS services to the internet or between AWS regions. Costs vary depending on the source and destination of the data transfer. It's crucial to understand these cost structures when planning your broadcasting infrastructure.

Strategies for Minimizing Data Transfer Costs

- **Region Selection:** Place your resources in the AWS region closest to your primary viewers or the majority of your operational crew. This reduces the distance data must travel, potentially lowering costs and improving latency. However, be aware of the pricing differences between regions and choose accordingly.



- **Optimizing Preview Signals for Operational Crew:** For live event productions, providing a preview signal to the operational crew is essential. Optimize these streams by using multi-viewer and proxy formats, which consume less bandwidth than full-resolution feeds. Adjusting the resolution and bitrate according to the needs of the preview can contribute to cost savings.
- **Efficient Use of NICE DCV for Operator Screens:** When leveraging remote desktop technologies like NICE DCV for remote access by operators, consider the number of screens and the data transfer involved. Optimizing the setup to balance performance and quality with the data transfer volume can lead to cost reductions. This might involve adjusting the screen resolution, frame rate, or utilizing compression settings where possible without significantly impacting the operational effectiveness.
- **Content Delivery with Amazon CloudFront:** Amazon CloudFront is a global content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to viewers with low latency and high transfer speeds. Utilizing CloudFront can significantly reduce data transfer costs by caching content at edge locations closer to your viewers, decreasing the need for data to travel across regions.

9.6 Tools for Cost Optimization

AWS provides essential cost management tools designed to help you monitor, analyze, and optimize your cloud spending. This compact guide highlights how to leverage these tools effectively, focusing on the unique needs of live cloud production.

AWS Cost Explorer

AWS Cost Explorer is a powerful analytics tool for visualizing and understanding your AWS spending. It is invaluable for tracking the costs associated with TriCaster Vector deployments, enabling you to:

- Analyze historical spending to identify trends and make informed budgeting decisions.
- Forecast future spending to ensure alignment with your financial plans.
- Drill down into specific cost drivers, such as instance types or regions, to pinpoint optimization opportunities.

AWS Budgets

AWS Budgets empowers you to set custom budget thresholds and receive alerts when your spending exceeds these limits. It's crucial for managing the costs of live production resources, allowing you to:

- Define budgets for TriCaster Vector instances and related AWS services.



- Receive real-time alerts to prevent overspending, particularly valuable for resources that may be left running unintentionally—a common scenario in dynamic production environments.

AWS Trusted Advisor

AWS Trusted Advisor offers personalized recommendations to optimize your AWS environment for cost, performance, and security. For TriCaster Vector deployments, it can help by:

- Identifying underutilized or idle resources, enabling you to scale down or terminate them to save costs.
- Recommending Reserved Instances or Savings Plans based on your usage, offering substantial cost savings over time.

Resource Tagging

Effective resource tagging is essential for granular cost tracking and optimization in live cloud productions. By tagging resources related to TriCaster Vector deployments (e.g., by project or department), you can:

- Allocate costs accurately and streamline cost reporting.
- Implement automated policies for resource management, enhancing cost efficiency.

Implementing Alerts and Budgets

To maintain control over your live production costs, it's important to:

1. Establish clear metrics that reflect your cloud spending and operational efficiency.
2. Utilize AWS Budgets to create and manage budgets tailored to your live production needs, setting up alerts for potential overruns.
3. Regularly review AWS Trusted Advisor recommendations to optimize your deployment for cost savings without impacting production quality.

9.7 AWS Calculator

Creating an accurate AWS cost estimate for a small to medium Live Cloud Event Broadcasting solution involves several components and configurations. This example aims to provide an overview of potential costs for 1 hour of production in the AWS US East (N. Virginia) region, focusing on the essential elements such as egress for remote desktop, monitoring feed egress, and the recommended EC2 and EBS configurations.

Cost Calculation Components



1. **Egress for Remote Desktop:**

- **Usage:** Utilization by operators for remote control and viewing, based on 1920x1080 60P NDlHX H.264 video quality.

2. **Egress for Monitoring Feed:**

- **Usage:** A multiview stream for monitoring purposes, based on the same video quality as the remote desktop usage.

3. **EC2 and EBS Configuration:**

- **EC2:** Recommended instance type for small to medium production needs.
- **EBS:** Recommended volume size and type to support the required video storage and playback for 1 hour of production.

Important Reminders for Customers

- **Cost Responsibility:** Customers are fully responsible for the costs incurred from using deployed solutions within their AWS account/environment. It's crucial to monitor usage closely and optimize resources to manage expenses effectively.
- **Estimation Accuracy:** The AWS Pricing Calculator offers an estimate of AWS fees without including taxes or potential discounts. Actual fees may vary based on factors such as actual usage patterns, changes in AWS pricing, or the application of reserved instances or savings plans.
- **Usage Factors:** The actual fees depend on various factors, including the specific configurations chosen, the duration of usage, and any additional AWS services employed during production.



Summary		Region	Upfront		Hourly								
		US East (N. Virginia)	\$	\$	\$	\$							
Group	Application / Services	Services	Upfront	Hourly	Configuration Summary								
Live Event Broadcasting > Small Size 1 - N. Virginia > Live Event Broadcasting	Vision Mixer - TriCaster Vector	Amazon EC2	0	\$ 2.84	Tenacity (Shared instances), Operating system (Windows Server), Workload (Monthly, Baseline: 0, Peak: 1, Duration of peak: 0 Day 1 Hr 0 Min), Advance EC2 instance (g4dn.4xlarge), Pricing strategy (On-Demand), Enable monitoring (disabled), EBS Storage amount (770 GB), DT Inbound: Not selected (0 TB per month), DT Outbound: Internet (10 GB per month), DT intra-Region: (0 TB per month)								
					Operating system	Advance EC2 instance	Workload	Duration of peak	Workload days	DT Inbound	DT Intra-Region	Storage amount	
					Windows Server	g4dn.4xlarge	Monthly	0 Day 1 Hr 0 Min				770 GB	
					Tenacity	Pricing strategy	Baseline	Peak	0	DT Outbound: Internet	DT Outbound: Not selected		
					Shared instances	On-Demand	0	1		10 GB per month			
Live Event Broadcasting > Small Size 1 - N. Virginia > Live Event Broadcasting	Replay System - Viz 3Play 3PV	Amazon EC2	0	\$ 2.39	Tenacity (Shared instances), Operating system (Windows Server), Workload (Monthly, Baseline: 0, Peak: 1, Duration of peak: 0 Day 1 Hr 0 Min), Advance EC2 instance (g4dn.4xlarge), Pricing strategy (On-Demand), Enable monitoring (disabled), EBS Storage amount (770 GB), DT Inbound: Not selected (0 TB per month), DT Outbound: Internet (5 GB per month), DT intra-Region: (0 TB per month)								
					Operating system	Advance EC2 instance	Workload	Duration of peak	Workload days	DT Inbound	DT Intra-Region	Storage amount	
					Windows Server	g4dn.4xlarge	Monthly	0 Day 1 Hr 0 Min				770 GB	
					Tenacity	Pricing strategy	Baseline	Peak	0	DT Outbound: Internet	DT Outbound: Not selected	Shared instances	
					Shared instances	On-Demand	0	1		5 GB per month			
Live Event Broadcasting > Small Size 1 - N. Virginia > Live Event Broadcasting	Graphics - Viz Engine	Amazon EC2	0	\$ 1.12	Tenacity (Shared instances), Operating system (Windows Server), Workload (Monthly, Baseline: 0, Peak: 1, Duration of peak: 0 Day 1 Hr 0 Min), Advance EC2 instance (g4dn.2xlarge), Pricing strategy (On-Demand), Enable monitoring (disabled), EBS Storage amount (100 GB), DT Inbound: Not selected (0 TB per month), DT Outbound: Internet (5 GB per month), DT intra-Region: (0 TB per month)								
					Operating system	Advance EC2 instance	Workload	Duration of peak	Workload days	DT Inbound	DT Intra-Region	Storage amount	
					Windows Server	g4dn.2xlarge	Monthly	0 Day 1 Hr 0 Min				100 GB	
					Tenacity	Pricing strategy	Baseline	Peak	0	DT Outbound: Internet	DT Outbound: Not selected	Shared instances	
					Shared instances	On-Demand	0	1		0 TB per month			
Live Event Broadcasting > Small Size 1 - N. Virginia > Live Event Broadcasting	Graphics Control Application - Viz Trio	Amazon EC2	0	\$ 1.57	Tenacity (Shared instances), Operating system (Windows Server), Workload (Monthly, Baseline: 0, Peak: 1, Duration of peak: 0 Day 1 Hr 0 Min), Advance EC2 instance (g4dn.2xlarge), Pricing strategy (On-Demand), Enable monitoring (disabled), EBS Storage amount (30 GB), DT Inbound: Not selected (0 TB per month), DT Outbound: Internet (5 GB per month), DT intra-Region: (0 TB per month)								
					Operating system	Advance EC2 instance	Workload	Duration of peak	Workload days	DT Inbound	DT Intra-Region	Storage amount	
					Windows Server	g4dn.2xlarge	Monthly	0 Day 1 Hr 0 Min				30 GB	
					Tenacity	Pricing strategy	Baseline	Peak	0	DT Outbound: Internet	DT Outbound: Not selected	Shared instances	
					Shared instances	On-Demand	0	1		5 GB per month			
Live Event Broadcasting > Small Size 1 - N. Virginia > Live Event Broadcasting	Graphics Control Application - Viz Trio	Amazon EC2	0	\$ 1.57	Tenacity (Shared instances), Operating system (Windows Server), Workload (Monthly, Baseline: 0, Peak: 1, Duration of peak: 0 Day 1 Hr 0 Min), Advance EC2 instance (g4dn.2xlarge), Pricing strategy (On-Demand), Enable monitoring (disabled), EBS Storage amount (30 GB), DT Inbound: Not selected (0 TB per month), DT Outbound: Internet (5 GB per month), DT intra-Region: (0 TB per month)								
					Operating system	Advance EC2 instance	Workload	Duration of peak	Workload days	DT Inbound	DT Intra-Region	Storage amount	
					Windows Server	g4dn.2xlarge	Monthly	0 Day 1 Hr 0 Min				30 GB	
					Tenacity	Pricing strategy	Baseline	Peak	0	DT Outbound: Internet	DT Outbound: Not selected	Shared instances	
					Shared instances	On-Demand	0	1		5 GB per month			
Live Event Broadcasting > Small Size 1 - N. Virginia > Live Event Broadcasting	Contribution - NDI Bridge	Amazon EC2	0	\$ 1.57	Tenacity (Shared instances), Operating system (Windows Server), Workload (Monthly, Baseline: 0, Peak: 1, Duration of peak: 0 Day 1 Hr 0 Min), Advance EC2 instance (g4dn.2xlarge), Pricing strategy (On-Demand), Enable monitoring (disabled), EBS Storage amount (30 GB), DT Inbound: Not selected (0 TB per month), DT Outbound: Internet (5 GB per month), DT intra-Region: (0 TB per month)								
					Operating system	Advance EC2 instance	Workload	Duration of peak	Workload days	DT Inbound	DT Intra-Region	Storage amount	
					Windows Server	g4dn.2xlarge	Monthly	0 Day 1 Hr 0 Min				30 GB	
					Tenacity	Pricing strategy	Baseline	Peak	0	DT Outbound: Internet	DT Outbound: Not selected	Shared instances	
					Shared instances	On-Demand	0	1		5 GB per month			
Live Event Broadcasting > Small Size 1 - N. Virginia > Live Event Broadcasting	Discovery Services - NDI Discovery	Amazon EC2	0	\$ 0.06	Tenacity (Shared instances), Operating system (Windows Server), Workload (Monthly, Baseline: 0, Peak: 1, Duration of peak: 0 Day 1 Hr 0 Min), Advance EC2 instance (t3.medium), Pricing strategy (On-Demand), Enable monitoring (disabled), EBS Storage amount (30 GB), DT Inbound: Not selected (0 TB per month), DT Outbound: Internet (5 GB per month), DT intra-Region: (0 TB per month)								
					Operating system	Advance EC2 instance	Workload	Duration of peak	Workload days	DT Inbound	DT Intra-Region	Storage amount	
					Windows Server	t3.medium	Monthly	0 Day 1 Hr 0 Min				30 GB	
					Tenacity	Pricing strategy	Baseline	Peak	0	DT Outbound: Internet	DT Outbound: Not selected	Shared instances	
					Shared instances	On-Demand	0	1		0 TB per month			
Live Event Broadcasting > Small Size 1 - N. Virginia > Live Event Broadcasting	Public IPs	Amazon Elastic IP	0	\$ -	Number of EC2 instances (6), Number of EIPs per instance (1)								
Live Event Broadcasting > Small Size 1 - N. Virginia > Live Event Broadcasting	VPC Live Event Broadcasting	VPN Connection	0	\$ -	Working days per month (22)								
Live Event Broadcasting > Small Size 1 - N. Virginia > Backing Services	Viz License Server	Amazon EC2	0	\$ 0.04	Tenacity (Shared instances), Operating system (Windows Server), Workload (Monthly, Baseline: 0, Peak: 1, Duration of peak: 0 Day 1 Hr 0 Min), Advance EC2 instance (t3.small), Pricing strategy (On-Demand), Enable monitoring (disabled), EBS Storage amount (50 GB), DT Inbound: Not selected (0 TB per month), DT Outbound: Internet (5 GB per month), DT intra-Region: (0 TB per month)								
					Operating system	Advance EC2 instance	Workload	Duration of peak	Workload days	DT Inbound	DT Intra-Region	Storage amount	
					Windows Server	t3.small	Monthly	0 Day 1 Hr 0 Min				50 GB	
					Tenacity	Pricing strategy	Baseline	Peak	0	DT Outbound: Internet	DT Outbound: Not selected	Shared instances	
					Shared instances	On-Demand	0	1		0 TB per month			
Live Event Broadcasting > Small Size 1 - N. Virginia > Backing Services	Viz Graphic Hub	Amazon EC2	0	\$ 0.11	Tenacity (Shared instances), Operating system (Windows Server), Workload (Monthly, Baseline: 0, Peak: 1, Duration of peak: 0 Day 1 Hr 0 Min), Advance EC2 instance (t3.large), Pricing strategy (On-Demand), Enable monitoring (disabled), EBS Storage amount (100 GB), DT Inbound: Not selected (0 TB per month), DT Outbound: Internet (5 GB per month), DT intra-Region: (0 TB per month)								
					Operating system	Advance EC2 instance	Workload	Duration of peak	Workload days	DT Inbound	DT Intra-Region	Storage amount	
					Windows Server	t3.large	Monthly	0 Day 1 Hr 0 Min				100 GB	
					Tenacity	Pricing strategy	Baseline	Peak	0	DT Outbound: Internet	DT Outbound: Not selected	Shared instances	
					Shared instances	On-Demand	0	1		0 TB per month			
Live Event Broadcasting > Small Size 1 - N. Virginia > Backing Services	Public IPs	Amazon Elastic IP	0	\$ -	Number of EC2 instances (2), Number of EIPs per instance (1)								
Live Event Broadcasting > Small Size 1 - N. Virginia > Backing Services	VPC Backing Services	VPN Connection	0	\$ -	Working days per month (22)								

9.8 Vizrt Flexible License Program

Understanding the cost model and licensing options for TriCaster Vector is essential for organizations looking to leverage this powerful live production tool within their cloud-based broadcasting workflows. TriCaster Vector operates on a Bring Your Own License (BYOL) model, allowing clients to tailor their licensing according to their specific needs and operational tempo.

Clients interested in acquiring a TriCaster Vector license have two primary avenues: direct engagement with Vizrt Sales personnel or through the extensive Vizrt Partner Channel. These channels offer personalized assistance to help clients select the most appropriate licensing options based on their production scale, feature requirements, and budgetary considerations.

Furthermore, TriCaster Vector is available under the Vizrt Flexible Access program, which introduces a high degree of flexibility and accessibility for users. This program offers subscription-based licensing with various commitment terms, ranging from a daily license for short-term projects and events, extending to multiple months for ongoing operations. This



flexibility ensures that organizations of all sizes and scopes can access TriCaster Vector's capabilities in a manner that aligns with their project timelines and budget constraints.

The BYOL model, combined with the Flexible Access program, positions TriCaster Vector as a versatile solution for live cloud production, accommodating the dynamic nature of live broadcasting environments. By offering different licensing terms, Vizrt ensures that clients can efficiently manage their licensing costs while enjoying the robust features and reliability of TriCaster Vector for their live production needs.



10 Integration with AWS Services

10.1 Integration Points with AWS Services

Throughout our guide, we've explored various integration points between TriCaster Vector and AWS services, highlighting how broadcasters can utilize AWS's robust cloud infrastructure to enhance live cloud production capabilities.

Computing

- **EC2 Instances:** TriCaster Vector benefits from the computational power of EC2 instances, especially those optimized for high-performance and graphic-intensive tasks. While traditional auto-scaling is not applicable, broadcasters can manually scale their EC2 instances to meet the demands of live production.

Storage and Content Delivery

- **Amazon EBS and S3:** For storing and managing video content, Amazon EBS provides high-performance block storage, while Amazon S3 offers scalable object storage. Both services ensure that data is readily available and durably stored.
- **Amazon CloudFront:** This CDN service integrates with TriCaster Vector to deliver content globally with low latency and high transfer speeds, ensuring an excellent viewer experience.

Networking and Connectivity

- **Amazon VPC:** A foundational integration point, Amazon VPC enables broadcasters to provision a logically isolated section of the AWS Cloud. Here, they can define and manage network configurations for improved security and network performance.
- **AWS Direct Connect:** For reduced latency and increased bandwidth throughput, AWS Direct Connect establishes a dedicated network connection from the broadcaster's premises to AWS.

Management and Monitoring

- **Amazon CloudWatch:** A critical service for monitoring the performance and health of AWS resources, CloudWatch allows broadcasters to track metrics, set alarms, and automatically react to changes in their TriCaster Vector environment.
- **AWS CloudTrail:** CloudTrail is essential for governance, compliance, and risk management within AWS accounts. By integrating with TriCaster Vector, broadcasters can log, continuously monitor, and retain a detailed event history of



actions taken across their AWS infrastructure, providing a clear audit trail of changes and aiding in operational analysis.

- **AWS CloudFormation:** For managing infrastructure as code, AWS CloudFormation enables the automated setup and provisioning of AWS resources, making the deployment of TriCaster Vector environments repeatable and consistent.

Security and Compliance

- **AWS Identity and Access Management (IAM):** IAM roles and policies are central to managing access to AWS resources securely, ensuring that only authorized users and services can interact with the TriCaster Vector deployment.
- **AWS Key Management Service (KMS):** KMS integrates with TriCaster Vector to manage encryption keys, allowing broadcasters to protect their stored content and comply with industry regulations.
- **AWS Secrets Manager:** Managing credentials and secrets securely is critical for any live production environment. AWS Secrets Manager helps broadcasters to securely encrypt, store, and retrieve credentials, such as database passwords or API keys, necessary for TriCaster Vector deployments. Integrating with Secrets Manager ensures that sensitive information is handled according to best security practices, reducing the risk of unauthorized access.

10.2 Leveraging AWS for Enhanced Capabilities

Integrating TriCaster Vector with AWS services can vastly enhance the capabilities of live cloud productions, offering broadcasters a range of powerful tools for streamlining workflows and expanding functionality. AWS Media Services, in particular, offers specialized services for creating and delivering video content efficiently. Here's how to integrate TriCaster Vector with AWS services for enhanced broadcasting capabilities.

- **AWS Elemental MediaLive:** Integrate with AWS Elemental MediaLive to encode live video streams for broadcast and streaming to multiple platforms. MediaLive can handle high-quality streams and is highly scalable, making it suitable for events of any size.
- **AWS Elemental MediaConnect:** Use MediaConnect to securely transfer live video feeds over the AWS network. It facilitates the sharing of live video between TriCaster Vector and other locations or services, providing flexibility and reliability for high-bandwidth video streams.
- **AWS Elemental MediaPackage:** Utilize MediaPackage to package live streams for delivery to different devices, customizing the video content for various platforms while maintaining content protection.




- **AWS Elemental MediaStore:** MediaStore serves as a video origination and storage service, offering the speed and consistency required for live streaming combined with the security and durability of AWS.

Leveraging AWS for Enhanced Capabilities

- **Scalability and Reliability:** AWS Media Services are designed to automatically scale with the number of viewers and the bandwidth required for your stream, ensuring consistent performance even during demand spikes.
- **Global Distribution:** AWS's extensive global network allows for the distribution of live streams with low latency, no matter where the audience is located.
- **Security and Compliance:** AWS services include built-in security features to protect your streams and ensure compliance with industry regulations. MediaPackage, for example, includes digital rights management (DRM) capabilities.

10.3 Example Use Cases and Scenarios

- **Large-Scale Live Events:** For events with a significant number of viewers, such as sports events or concerts, integrating TriCaster Vector with AWS Elemental MediaLive and MediaConnect ensures that live feeds are reliably encoded and distributed across various platforms and devices.
- **24/7 Live Channels:** For continuous live broadcasting, AWS Elemental MediaPackage can be used to create channel-like streams that are packaged on-the-fly and delivered to viewers without interruption.
- **Multiregional Broadcasts:** When streaming content to a global audience, AWS Elemental MediaStore can act as the origin store, providing viewers with low-latency access to live streams.
- **Secure Content Delivery:** In scenarios where content protection is paramount, AWS Media Services offer encryption and DRM to ensure that only authorized viewers can access the streams.



11 Quality Assurance and Testing

Quality Assurance (QA) and testing are critical components in ensuring the success and reliability of TriCaster Vector deployments for live cloud productions. This chapter outlines strategies for testing and best practices for QA, enabling broadcasters to confidently deploy and manage their live production environments.

11.1 Testing Strategies

When preparing to test TriCaster Vector deployments, it's crucial to simulate real-world scenarios that reflect your specific use cases.

1. **Define Use Case Scenarios:** Start by clearly defining the scenarios in terms of the number of inputs and outputs, transport streams, codecs, and quality. This should mirror the actual operational environment as closely as possible.
2. **Simulate Operational Load:** Include simulations of remote desktop load and all utilized features such as stream recording. This will give you a realistic picture of how the system performs under typical operational conditions.
3. **Metrics Collection:** Collect and analyze metrics during testing to make informed decisions. This data will be crucial in understanding system performance and identifying any potential bottlenecks or issues.

11.2 Testing and Production environments

It is highly recommended to have separate environments for testing and production.

The best practice in system administration advises isolating the various broadcast activities occurring concurrently in at least two environments:

- **Testing / QA:** This should be almost a copy of the production environment (maybe with much less resources) but overall same setup. This is used by the specific departments (QA, Support, group of advanced users) to do their own testing, including load testing. (one environment can be used for load testing and another for manual testing if needed, so that it does not disrupt all QA under heavy load). Automation testing can also be run here or have its own environment. Data generated here is random depending on the needed tests to cover all edge cases.
- **Production:** This is where value is created for customers and/or the business. Production is a highly visible environment that deeply impacts Grupo Globo operation. The production environment should be as free of errors as possible.




This best practice view emphasizes that all these activities should not occur in the same environment.

The key benefit of multiple environments is that content validation occurs in a non-production environment, thus ensuring that any modifications to the model logic result in accurate figures for existing content. It also allows for performance tests on new code to prevent any reports and subsequent analysis from grinding to a halt.

11.3 Performance and Stress Testing

To stress test your TriCaster Vector deployment, consider creating a specific VPC designed to simulate video capture and transcoding.

1. **Simulated Stream Provisioning:** Configure the VPC to support a variety of NDI or SRT streams, with distinct bit rates if possible, originating from both on-premises setups and internet sources. These streams will be utilized in TriCaster Vector to facilitate diverse output compositions, ensuring a versatile and robust production capability. This setup allows for the seamless integration of inputs with varying bit rates, enhancing the flexibility and quality of the broadcast output..
2. **Activate Features:** Ensure comprehensive activation and testing of all TriCaster Vector features tailored for specific use cases. This encompasses a wide range of functionalities, including graphics overlays, transitions, multi-viewer outputs, and extends to macros automation for streamlining operations, integration with multiple sources for diverse input, advanced streaming capabilities for robust content delivery, format conversion to ensure compatibility across different platforms, and export options for distributing content. Each of these elements should be meticulously verified to guarantee a seamless production experience.
3. **Metrics Monitoring:** As per the guidance in the "Performance Efficiency" chapter, continuously monitor and record performance metrics. This will help in assessing the system's ability to handle the defined load.



11.4 Quality Assurance Best Practices

1. **End-to-End Testing:** Regularly schedule and conduct end-to-end tests of the entire live production pipeline. This ensures all components work harmoniously and meet the expected quality standards.
2. **Iterative Testing:** Implement an iterative approach to testing, where feedback from each test phase is used to make improvements. This helps in progressively refining the system's performance and reliability.
3. **User Experience Focus:** Pay special attention to the end-user experience during testing. This includes assessing the quality of the video and audio output, latency, and the overall viewer experience.
4. **Automation:** Utilize TriCaster Vector's macro automation capabilities to simulate repetitive tasks and operator actions for in-depth performance analysis. This approach enables the systematic evaluation of operational efficiency and the identification of potential areas for improvement, ensuring a more streamlined and effective production workflow.



5. **Documentation:** Keep detailed records of all test scenarios, methodologies, and results. This documentation is invaluable for future reference and for guiding ongoing improvements.

Exploring Chaos Engineering

1. **Fault Injection:** Embrace chaos engineering techniques by introducing faults into the system, including simulating 100% CPU usage, network disruptions, latency issues, and simulating problems with video feed signals. This strategy aids in comprehensively understanding the system's behavior under various stress conditions, facilitating targeted enhancements in its resilience and performance.
2. **Resilience Planning:** Use the insights gained from chaos engineering to enhance the resilience of your deployment. This could involve optimizing resource allocation, enhancing failover mechanisms, or improving system monitoring.



12 Documentation and Support

Ensuring that our customers have access to comprehensive documentation and robust support is a cornerstone of our service commitment at Vizrt. This chapter outlines the resources available for user manuals and technical documentation, details on support channels and resources, and information about our community and forums.

12.1 User Manuals and Technical Documentation

Vizrt provides extensive documentation to assist our users in maximizing the benefits of our solutions. Our official documentation, including user manuals, technical guides, and product release notes, is accessible online. These resources are designed to offer in-depth insights into product functionalities, setup processes, and operational best practices.

- Access the **Vizrt documentation** online at: <https://documentation.vizrt.com/>

12.2 Support Channels and Resources

Vizrt's global support organization is structured to deliver expert assistance around the clock, in over 30 languages worldwide. Our centralized support infrastructure ensures that all Vizrt customers receive unified and standardized service, regardless of their location.

How to Receive Support:

- **Vizrt Global Support:** Our knowledge, workflows, and procedures form the backbone of our regional support teams, ensuring coordinated operations for a seamless support experience. Vizrt Support offers media companies peace of mind through:
 - 24/7 access to our experts worldwide
 - Access to the latest software versions
 - Entry to the Vizrt Support Portal and Knowledge Base
 - A dedicated support team for issue resolution
- **Vizrt Support Portal:** Provides access to our services and a platform for interaction with Vizrt Global Support. Visit the support portal at <https://www.vizrt.com/support>.

Technical Support Tiers:

The Vizrt Global Support Handbook outlines our Software Maintenance and Support Programs, detailing the scope and conditions of technical support available to our customers.



This handbook is integral to the Vizrt License and Services Agreement and defines the entire agreement regarding software maintenance and support.

- The handbook, including details on support tiers and SLAs, is available at:
<https://docs.vizrt.com/Vizrt-Global-Support-Handbook.pdf>
- **Service Level Agreement (SLA):** Section 2.3 of the Global Support Handbook specifies the services and response commitments, granting customers access to Vizrt resources and knowledge repositories. Full access to Viz University is included at all support levels.

For additional resources, please refer to the section 15.1 Additional Resources and References in the Appendices. This section offers a wealth of information and links to further resources essential for enhancing your deployment and operational proficiency.

13 Versioning and Updates

In the rapidly evolving landscape of live production technology, staying current with the latest software versions and updates is crucial for maintaining system performance, security, and access to new features. This chapter outlines the approach to versioning and updates for Vizrt products, including TriCaster Vectar, ensuring that customers can efficiently manage their systems and capitalize on advancements in technology.

13.1 Maintaining Your Environment Secure

Ensuring the security of your live production environment is crucial. This involves the regular rotation of programmatic system credentials and cryptographic keys to safeguard against unauthorized access and potential security breaches. Follow these steps to maintain the integrity and security of your system:

1. **Identify Credentials and Keys:** Start by identifying all programmatic credentials and cryptographic keys used within your environment.
2. **Rotate Credentials:** Regularly rotate these credentials according to a defined schedule or in response to specific security events.
3. **Update Systems:** Ensure all systems utilizing these credentials and keys are updated accordingly to prevent disruptions.
4. **Audit and Monitor:** Regularly audit access logs and monitor for unauthorized access attempts, adjusting security measures as needed.

13.2 Managing Solution Updates

Staying Updated:

Keeping your Vizrt products updated is vital for accessing the latest features and ensuring the highest level of security and performance.

- **Updates Notification:** Users are informed about available updates through the Vizrt public website and via email, providing immediate awareness of new releases.
- **Manual Check:** For additional updates, users can manually check the Vizrt FTP at <https://download.vizrt.com/>.

Update Process:

Updating Vizrt products involves a straightforward process designed to minimize disruption:



1. **Backup:** Always back up configurations and data before starting the update process.
2. **Review Release Notes:** Examine the release notes for update details, improvements, and instructions.
3. **Test in a Non-Production Environment:** Apply and test updates in a staging environment to ensure compatibility.
4. **Apply the Update:** Follow the provided instructions to complete the update process.
5. **Verify:** Confirm that all systems and new features are functioning correctly post-update.

Access to Updates:

- **Vizrt Support Portal and FTP:** Latest software versions and updates are available for registered users on the Vizrt Support Portal and FTP.
- **Software Maintenance and Support Program:** Enrolled customers have access to all updates, ensuring systems stay current with Vizrt's innovations.

13.3 Products License

Vizrt customers receive their product licenses during the purchase process, which may include temporary or definitive license keys for use in the Vizrt License server powered by WIBU.

- **Activation Process:** Customers are provided with an activation code or ticket by CodeMeter License Central, specifying the software's operational scope and offline usage terms.
- **Software Installation:** Upon first launch, users enter the ticket, prompting communication with CodeMeter License Central to validate and create a temporary offline license, ensuring the software operates seamlessly without a constant Internet connection.

13.4 Emergency Maintenance

When operating TriCaster Vector on AWS, it's crucial to have an emergency maintenance plan to handle fault conditions effectively. This section provides instructions on managing such situations, ensuring business continuity and minimal disruption to your live broadcast.

Handling Performance Issues During a Live Broadcast



1. **Prioritize Program Output and UI Responsiveness:** TriCaster Vector is designed to prioritize program output and respond to user interface inputs even when system resources are constrained.
2. **Reduce System Load:** If you encounter performance issues, consider reducing the system load. This could involve decreasing the number of active feeds or stopping the recording of feeds temporarily to allocate more resources to critical tasks.
3. **Monitor System Health:** Use AWS CloudWatch to monitor the system's health and performance in real-time. Look out for alerts that may indicate issues with your instances or services.
4. **Engage Failover Procedures:** If you've set up a failover mechanism as part of your business continuity plan, as discussed in the "Reliability and Resilience" chapter, engage these procedures to switch to standby resources if necessary.
5. **Post-Event Analysis:** After resolving the immediate issue, conduct a detailed analysis to determine the root cause. Review logs, metrics, and system reports to understand what happened and why.
6. **Adjust Architecture:** Based on your findings, make necessary adjustments to your architecture to enhance resilience and prevent similar issues in the future.

Dealing With Storage Connectivity Issues

1. **Initial Response:** If you lose connection with your storage, which could manifest as an inability to access content or session data, the first step is to attempt to close and reopen the TriCaster Vector software and session.
2. **Verify Storage Health:** Check the AWS Console to confirm that the storage (such as EBS or S3) is operational and in a healthy state. Review any alerts or notifications that may indicate problems with the storage service.
3. **Restart Services:** If storage is confirmed healthy but issues persist, consider restarting the affected services or instances, ensuring that you follow any set protocols for service restarts to avoid data corruption.
4. **Data Recovery Procedures:** If data access issues continue, follow your established data recovery procedures. This may involve restoring data from backups or engaging AWS support for assistance.
5. **Documentation and Reporting:** Document the incident and actions taken to resolve it. This documentation can be crucial for future reference and for improving your emergency response plans.

14 Conclusion and Next Steps

As we conclude this guide on optimizing TriCaster Vector deployments on AWS, it's important to consolidate the key points and outline actionable next steps to ensure deployment success. Below is a summary of essential considerations, along with recommendations and resources for ongoing support.

14.1 Summary of Key Points

- **Performance Efficiency:** Optimal selection and configuration of AWS resources are critical for the high-performance delivery of live cloud productions.
- **Reliability and Resilience:** A robust business continuity plan, including well-defined strategies for high availability and disaster recovery, ensures uninterrupted broadcasts.
- **Scalability:** Manual resource scaling, informed by continuous performance monitoring, is essential for addressing the dynamic demands of live production.
- **Security:** Implementing AWS security best practices protects your content and maintains compliance with industry regulations.
- **Integration:** Leveraging AWS services enhances the capabilities and efficiency of your TriCaster Vector deployments.

14.2 Recommendations for Deployment Success

1. **Continuous Learning:** Stay informed about the latest AWS features and best practices by regularly reviewing AWS documentation and engaging with Vizrt resources.
2. **Quality Assurance:** Conduct thorough testing and monitoring of your deployments to ensure they meet your production requirements and audience expectations.
3. **Iterative Improvement:** Apply an iterative approach to your operations, continuously refining your processes and infrastructure based on performance data and user feedback.
4. **Community Engagement:** Participate in forums and user groups to share experiences, gain insights, and stay connected with the broader Vizrt and AWS communities.



14.3 Contact Information for Further Assistance

- **Vizrt Global Support:** For specialized assistance with TriCaster Vector, reach out to Vizrt Global Support. Their expertise and support infrastructure are invaluable resources for resolving technical issues and optimizing your deployment [Vizrt Support](#).
- **AWS Support:** AWS offers a range of support plans to fit different needs, from development support to enterprise-level assistance.
 - **AWS Support Center:** AWS Support
 - **AWS Knowledge Center:** Knowledge Center

Next Steps

1. **Review and Implement:** Revisit the guide and implement the recommended strategies and configurations for your TriCaster Vector deployments.
2. **Training and Certification:** Consider enrolling in relevant AWS and Vizrt training courses to enhance your technical skills and earn certifications.
3. **Feedback Loop:** Establish a mechanism to gather feedback from your production team and audience to inform future improvements.

By following these guidelines and utilizing the available resources, you're well on your way to achieving a successful and efficient TriCaster Vector deployment on AWS. As you move forward, remember that the journey to operational excellence is continuous. Embrace innovation, seek support when needed, and always strive for the highest quality in your live cloud productions.



15 Appendices

The Appendices section serves as a comprehensive reference point at the conclusion of our guide on optimizing TriCaster Vector deployments on AWS. It is designed to provide you with easy access to supplementary information that will support your understanding and application of the content discussed throughout the guide.

15.1 Additional Resources and References

For those looking to deepen their knowledge and expertise in TriCaster Vector deployments and cloud-based live production workflows, a wealth of resources is available. The following list provides a starting point for further exploration and learning.

User Documentation

- **Vizrt Documentation:** The official [Vizrt Documentation](#) portal is a comprehensive resource for all product documentation related to TriCaster Vector and other Vizrt products. It is essential for users seeking detailed operational guidance.
- **TriCaster Vector User Manual:** The [user manual for TriCaster Vector](#) is an invaluable resource for users to learn how to perform common operations. It offers a balance of user-friendly guidance for everyday tasks and in-depth reference material for when detailed information is needed.

Learning and Certification

- **Viz University:** Offers a range of courses, including those focused on cloud adoption and TriCaster Vector operation. By completing these courses, users can gain certifications that demonstrate their proficiency and commitment to mastering live production technologies. Signing up at [Viz University](#) as a freelancer or media professional.
- **Vizrt Official YouTube Channel:** For visual learners, the [Vizrt Official YouTube Channel](#) provides a wealth of tutorial videos. These tutorials cover various aspects of setup, operation, and creative uses of Vizrt products, offering practical insights and tips.
- **Vizrt Video Tutorials:** For visual learners and those who prefer step-by-step video guidance, Vizrt's official Vimeo channel at [Vizrt Videos on Vimeo](#) offers a collection of tutorials and instructional content. These videos cover a range of topics from basic operations to advanced features of TriCaster Vector and other Vizrt products.



- **Community and Forums:** The Vizrt Community forum is a vibrant space for users to seek help, share knowledge, and connect with other professionals in the industry. This platform encourages the exchange of ideas, solutions to common challenges, and discussions on best practices. Join the conversation at the Vizrt Community forum: <https://forum.vizrt.com/>.

NDI Documentation

For comprehensive information on NDI technology, resources, and tools, please visit the official NDI documentation. These resources provide in-depth knowledge for leveraging NDI in your live production environments, covering everything from basic tools and setup guides to detailed technical white papers.

- **NDI Tools and Utilities:** Understand how to use and integrate NDI tools within your workflow. Visit [NDI Tools](#).
- **NDI Connected Community:** Space for seamless communication and exploration. Visit <https://community.ndi.video>.
- **NDI Guides:** Access a variety of guides that detail the operational aspects of NDI technologies. Check out [NDI Guides](#).
- **NDI White Paper:** Gain insights into the technical background and capabilities of NDI by reading the official white paper, "Introduction to NDI". Available at [NDI Introduction](#).

These resources are invaluable for both novice users and seasoned professionals looking to enhance their live production setups with NDI technology.

AWS Documentation and Best Practices

- **AWS Documentation:** AWS provides extensive documentation on all of its services and best practices at [AWS Documentation](#). This resource is regularly updated with the latest information on AWS features, services, and architectural recommendations.

15.2 Frequently Asked Questions

Q: What are the key considerations when selecting an AWS EC2 instance for TriCaster Vector?

A: When selecting an EC2 instance, consider the specific needs of your live production, such as the number of video inputs/outputs, resolution, frame rate, and expected workload. Evaluate the instance's computing power, memory, storage, and network capacity. AWS



provides a variety of instances, and for graphics-intensive tasks, GPU instances like the G4dn or G5 series are recommended.

Q: How can I ensure high availability for my live broadcasts on AWS?

A: To ensure high availability, deploy your TriCaster Vector across multiple Availability Zones within an AWS Region. Regularly back up your configuration and stream data to Amazon S3 and consider using Viz Now or AWS CloudFormation for quick re-deployment if needed.

Q: What are the best practices for securing my live production environment on AWS?

A: Secure your environment by following the principle of least privilege with AWS IAM, use security groups and network ACLs for fine-grained access control, enable encryption for data at rest using AWS KMS, and ensure all data in transit is encrypted. Regularly review AWS CloudTrail logs and utilize AWS Secrets Manager to manage sensitive credentials.

Q: How do I handle storage for high-resolution video content in AWS?

A: Use Amazon EBS for high-performance block storage, especially for content that requires frequent access. For long-term storage and archiving, Amazon S3 is more cost-effective. Implement lifecycle policies in S3 to transition older content to lower-cost storage classes like S3 Glacier.

Q: How can I monitor the performance of my TriCaster Vector live production?

A: Utilize Amazon CloudWatch to monitor metrics such as CPU utilization, network bandwidth, and disk I/O operations. Set up CloudWatch alarms to get alerted on any performance issues.

Q: Where can I find tutorials and additional learning resources for TriCaster Vector?

A: Official tutorials and learning resources can be found on the Vizrt documentation portal, Viz University for certification courses, and Vizrt's Vimeo channel for video tutorials. AWS also provides documentation and best practices guides that are invaluable when deploying live production environments.

Q: What should I do if I encounter performance issues during a live broadcast?

A: If performance issues arise, check Amazon CloudWatch for real-time metrics to identify the bottleneck. You may need to scale up your EC2 resources or adjust your network settings. Post-event, conduct a thorough analysis to identify the root cause and adjust your architecture to prevent future issues.