

Putting Broadcast-Quality Online Video
Workflows in the Cloud:
Yes You Can and Why You Should

F R O S T  S U L L I V A N

A Frost & Sullivan White Paper

Avni Rambhia, Industry Principal, Digital Media

Dan Rayburn, Industry Principal, Digital Media

Mukul Krishna, Senior Global Director, Digital Media

Executive Summary..... 3

Introduction 3

Why Should You Put Workflows in the Cloud?..... 4

Building an End to End Live Workflow in the Cloud 5

Production 6

Ingest and Acquisition 6

Encoding and Ad Insertion 7

Publishing and Delivery 8

Protection and Playback..... 8

Orchestration, KPI and Telemetry 9

Solution Brief: Microsoft Azure Media Services..... 9

The Bottom Line..... 11

EXECUTIVE SUMMARY

Over the top (OTT) video delivery is a complicated business, with device profiles and delivery technologies evolving rapidly and requiring specialized technical expertise. Quality of experience is critical to ensuring subscriber loyalty in the fickle and fragmented online environment. With business models and long-term brand relevance hinging more and more on the online video experience, media companies can no longer treat online streaming delivery as an afterthought. Smart infrastructure planning matters immensely in the online world, as achieving scalable, high-quality, cost-effective video on demand (VOD) over the Internet is tricky and easier said than done.

Global live event broadband-cast is a formidable technical challenge. From signal acquisition through production tasks, including commentary insertion and captioning, low-latency compression and protection to multi-platform delivery and fragmented device playback, many disparate components need to come together in perfect synchronization to ensure high-quality OTT delivery of a live event. Scalability and reliability are paramount, and achieving these objectives has traditionally necessitated expensive equipment outlays.

Capital investment requirements have been particularly challenging for short-duration special events where demand is hard to forecast, and broadcasters might experience a huge surge in users and a doubling or tripling of channel counts for short time durations of days or weeks. To fully monetize the potential of these live events online, broadcasters are increasingly turning to the cloud for its ability to deliver broadband-casts in a flexible, massively scalable and reliable way that's also less risky and more cost-effective than the traditional approach of building out on-premise infrastructure.

As a community, we have only just started to realize the tremendous value proposition the cloud offers in enabling a more streamlined, efficient, reliable, scalable, secure and cost-effective media workflow. Whether for live, live linear or on-demand content, the cloud is positioned to revolutionize media workflows over the next few years. In this paper, we'll show how many media workflow aspects are already being orchestrated through the cloud, and have achieved levels of maturity, scale and reliability that make it a viable approach for premium content owners and broadcasters. Cloud readiness has been fueled by a very vibrant vendor community, which we will discuss through the example of Azure Media Services, a cloud-based online video platform from Microsoft that has stepped up to the plate to enable professional media companies to exploit cloud infrastructure as never before.

INTRODUCTION

Cloud-based media workflows have truly come of age. Proof of this state of maturity is evidenced in the cloud-based, broadband-cast of the Sochi Winter Olympics by NBC Sports, where 41 live video feeds were delivered to tens of millions of connected devices worldwide over 17 days at bit rates ranging from 3.5 Mbps to 200 Kbps. Sochi was the first Olympics where online content delivery was achieved entirely in the cloud, from acquisition and encoding to packaging and streaming. This shows how far cloud-based technologies have come in fulfilling their promise of bringing cost-effective scalability, flexibility and agility to online video workflows, and serves as a harbinger of the future. It is clear that today, companies have progressed from experimenting in the cloud or simply executing final packaging and streaming through the cloud to actually putting entire mission-critical production workflows in the cloud.

As a result, TV Everywhere is moving from a trendy buzzword to ubiquitous reality. Overall for the Sochi Olympics, Microsoft Azure Media Services supported five broadcasters, including NBC, CBC, Viasat, América Móvil and NHK, across 22 countries in four continents with a population of 900 million; the magnitude of this achievement underscores the massive scale that the cloud can put at a broadcaster's fingertips today. The rest of this paper looks at the end-to-end component landscape and outlines key considerations for putting broadcast-quality video workflows in the cloud, using actual deployment details specifically from NBC's transmission of the Sochi Winter Olympics as a continuous case study in how the cloud can be effectively leveraged for production-grade online video workflows.

Why Should You Put Workflows in the Cloud?

Media companies know that to achieve maximum advertisement revenue and subscriber reach with ad-supported and subscription-based content, reaching every connected device with a solid quality of experience is critical. However, the cost and complexity of independently delivering large volumes of on-demand or live video to every target device form factor—accounting for variances in resolution support, format support, security support and playback environment support—are very high. The cloud is a far more effective alternative to on-premise deployments for TV everywhere workflows—an important consideration for media companies as they're being pressed to offer always-on, always-current content to a growing array of devices.

It is important to realize there are two distinct aspects to scalability: being able to scale the number of live channels and being able to scale the number of viewers per channel. The former is crucial for live events, while the latter is critical for all forms of VOD and live delivery.

Live events are a particularly compelling case study for the cloud. They are attractive to make available online, as studies show that they can drive nearly 10 times as much user engagement as on-demand video and because they offer straightforward monetization opportunities through broadcast advertisement revenue. The economics of a traditional "owned infrastructure" approach is hard to reconcile with live events. A media company preparing for a major sports event such as the Sochi Olympics would be required to spend massive amounts of money and time deploying infrastructure for two weeks of high-volume operations (including CAPEX, equipment rentals or purchase and deployment, and staffing), and then be prepared to dismantle this overhead once the event and its excitement are over. Often companies are forced to select only specific events to take online when they are upper-bounded by CAPEX considerations. In contrast, the flexible, just-in-time, pay-as-you-go nature of the cloud make it perfectly suited for scenarios that demand extreme elasticity and scalability – not just scalability in terms of concurrent viewers, but also scalability and elasticity in terms of concurrent channels. Thus, broadcasters can deliver significantly richer content choices (for example, broadcasting every event as NBC did at the Sochi Olympics) by bringing the viewer closer to an event through multiple camera angles; by enabling continuous broadcast of long-running events such as the ability to broadcast every hole in a golf match; or broadcasting every EPL game. Resources can not only be quickly spun up as necessary with very reasonable cost, but can also be spun down once the event concludes, which results in considerable long-term savings without compromising short-term revenue.

The cloud is also powerful at delivering high levels of reliability and redundancy in a cost-effective manner. Disaster recovery, for example, can be put into place by a cloud-based provider more easily, quickly and at lower cost than if a broadcaster attempted to do this on its own. Further, well-architected cloud solutions will provide multiple levels of redundancy at the data center level, channel level, and computing resource level. This is critical to ensuring reliable quality of experience for the end users and, ultimately, maximum monetization.

NBC Sports reported that 10.8 million hours of streamed video were consumed during the Sochi Winter Olympics – more than triple the hours of video streamed for the Vancouver Games. Of these, 80% were live streams consumed via NBCOlympics.com and NBC Sports Live Extra apps. Over the course of the event, over 2 million users accessed the content. Instead of acquiring, installing and maintaining racks of new encoders and servers to handle these feeds for the couple of weeks that the Olympics were on, and incurring the personnel costs to spin these resources up and down in reliable and redundant fashion, NBC Sports found a far more cost-effective and performance-effective option in the cloud. Cloud workflows also make sense for small events. The cost-savings that come from avoiding the need to deliver racks of hardware infrastructure on site for a one-time event are significant.

Another key advantage the cloud offers is economy of scale, with much more cost-effective pricing for storage, bandwidth, and worldwide client delivery as compared to the total cost of ownership for an on-premise alternative. Cloud vendors can amortize their equipment investments across a number of customers, and can afford far greater bandwidths and connections than a single company would be able to do. Furthermore, leading cloud platform vendors are constantly innovating and optimizing their own offerings and steadily pass on savings from those optimizations to their customers, which Frost & Sullivan believes will make cloud platforms even more affordable and compelling over time.

BUILDING AN END-TO-END LIVE WORKFLOW IN THE CLOUD

In a traditional on-premise architecture, components from many vendors come together to achieve an end-to-end workflow, and these have to integrate and work flawlessly together to achieve high-quality programming and content delivery. The same is true for the cloud as well. A typical workflow is depicted in the image below.

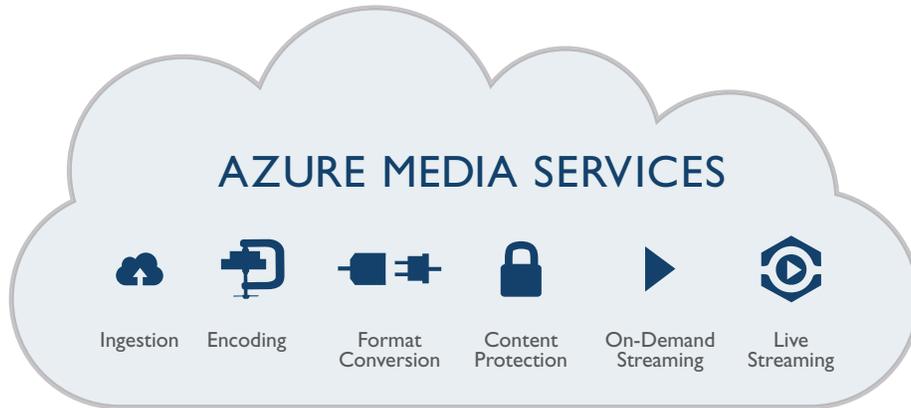


In a cloud-based media platform like Microsoft Azure Media Services, these components run on the actual cloud computing infrastructure and are connected externally through a CDN. On top of these functional components sit infrastructural components, including:

- the content management system (CMS),
- instrumentation to stitch the various components together,
- key performance indicator (KPI) monitoring, and
- failover and redundancy management, which may be spread over a single data center or – for larger operations – across multiple data centers.

In simpler workflows, a single cloud vendor may work to deliver all these components. However, for production-grade workflows, many individual components from best-of-breed partners need to be stitched together in a programmatic and customized fashion to deliver a workflow that works best for a given client and a given scenario. For its Sochi Winter Olympics programming, NBC Sports chose iStreamPlanet’s Aventus for ingest, encoding and API-driven ad insertion. Adobe enabled playback and the analytics, as well as the TV Everywhere authentication, which was used in lieu of full-fledged DRM. Akamai provided CDN services for OTT video

streaming. All of this was implemented on Microsoft Azure Media Services, a state-of-the-art, cloud-based media platform that combines the power and geographical reach of the Azure platform with decades of media engineering experience of the Microsoft Media team. CMS integration and instrumentation was implemented by Aventus by leveraging APIs provided by Azure Media Services on one hand and integrating with NBC's own content management systems on the other. The workflows were redundantly implemented across two US Azure data centers – East and West – with additional per-channel redundancy built into each workflow. A generic overview of the component functionality supported within Azure Media Services is depicted below.



In this section, we take a deeper look at each component of a cloud-based workflow and discuss factors to be considered in a cloud-based implementation of that component. The Sochi workflow will be used as a real-world showcase of how these components can be stitched together in the cloud, providing live event broadcasters flexibility and agility wrapped around stringent SLAs.

Production

The production process begins at camera feed acquisition and includes major production tasks such as contribution encoding as well as numerous minor tasks, including generation of highlights, sub-clipping, captioning, translations, metadata insertion, camera angle selection (where applicable), and more. These tasks are typically executed at the broadcaster's facility. For remote events, official camera feeds are typically backhauled to a central production facility using dedicated satellite or IP links. Once the final programming is available, it is typically compressed into an MPEG-2 transport stream format and is now ready for ingest into the cloud workflow for online content delivery.

For Sochi, NBC acquired 26 official live feeds directly from the Sochi Olympics Broadcasting Services (OBS) through the International Broadcast Center in Sochi, and added another 15 channels locally at its central facility in Stamford, Conn. All workflow management, advertisement operations and highlight generation tasks were implemented there as well. Together, these feeds were delivered via a dedicated IP link to iStreamPlanet at the Switch SuperNAP in Las Vegas, Nev., for aggregation and ingest into the cloud.

Ingest and Acquisition

For live content, ingest is the process by which live feeds are delivered into the cloud-based media processing system. For file-based scenarios, content is uploaded into the cloud service from the content owner's media asset management system or content management system, typically using a managed file transfer layer. With

a cloud-based architecture, the ideal situation is that the cloud service provider has a dedicated link to the broadcaster or programmer; this helps ensure minimum latency and maximum reliability. For added redundancy, content may be ingested into more than one data center run by the cloud service provider.

For Sochi, there were multiple levels of redundancy built into the ingest component. For example, content was acquired via IP over fiber networks but for key events, satellite backup was kept ready. Having received the NBC feeds, iStreamPlanet replicated these streams at the Switch SuperNAP and sent them to Azure's US East and West data centers via private IP paths for onward processing in the Aventus platform running within Azure Media Services.

Encoding and Ad Insertion

Encoding or transcoding of the audio, and particularly the video streams, is the most computationally intensive and technologically complex aspect of any workflow.

From a compression perspective, most video encoding will be done in AVC and most audio will be compressed using AAC. From a device reach perspective, anywhere between four and 15 video profiles are generally generated per input stream, spanning different resolutions and bit rates. During encoding, ad markers are also inserted into streams in conjunction with the underlying CMS. Automated quality monitoring tools are instrumented into this component, typically, to ensure glitch-free input and output and to spin-up backup resources where necessary.

Because of its complexity, it is crucially important that this component be implemented in a cloud-aware manner to provide maximum performance and error-resiliency. For example, the encoding will be running in a virtualized fashion, but there should be redundant VMs running per stream such that if one VM fails, another can smoothly take over without an interruption in the output. Similarly, if encoding workloads rise, the instrumentation layer should automatically spin up additional VMs to handle the additional loads.

For Sochi, proprietary encoding technology within Aventus was used, executing on Azure Media Services. Audio was encoded as 56 Kbps stereo HE-AAC v1, while video was encoded as AVC in seven separate profiles with bit rates ranging from 200 Kbps to 3.5 Mbps. iStreamPlanet worked with NBC Sports to create a custom CMS for the Olympics to perform ad insertion and slate insertion. During encoding, ad markers were inserted into the live streams via APIs to enable dynamic mid-roll ad insertion at a downstream point of time. Specifically, Aventus received REST-based API calls for ad insertion from ad operators via its Video CMS, which it converted to XML-formatted SCTE-35 cue messages and inserted into a text track. Azure Media Services then took the SCTE-35 messages and transformed them into the manifest types needed for each device type. Closed captioning was included in the feeds and was passed through the system for downstream use.

Aventus is built for the cloud from the ground up, with in-built protection against failure and automatic scalability features. Azure Media Services itself provided additional layers of redundancy. At the first level, the entire workflow was replicated across the US East and West data centers; at the second level, within each data center, individual streams were being redundantly encoded. In the case of any data glitches, Azure Media Services ensured that clean data was mirrored from one data center to the other, so that each center ultimately had pristine and complete libraries of encoded content.

Publishing and Delivery

Online video streaming to devices today is predominantly done using HLS and to computers is typically done using HDS or Smooth Streaming. DASH is another key emerging standard for adaptive bit rate (ABR) video streaming. During cloud-based media processing, the pre-created audio and video compressed streams are packaged into these different formats, with appropriate translations necessary for closed captioning, metadata and ad insertion markers. The finished packages are then immediately passed on to the CDN for live delivery and adjusted as necessary for later VOD use (for example, opening slates may be shortened), and then saved with manifests into cloud-based storage servers.

Both Aventus and Azure Media Services are capable of delivering live video in multiple HTTP-based adaptive streaming formats such as Adobe's new ABR format HDS, Apple's HLS and DASH. In the case of Sochi, Aventus handed off mp4 fragments to Azure Media Services. Following the handoff, Azure Media Services used its Dynamic Packaging feature to dynamically re-multiplex streams to HLS and HDS formats for delivery to end users. This included translation of the ad markers inserted by Aventus during the encoding process into HDS- and HLS-compatible markers, for downstream video players to interpret and act on appropriately. The live content was pulled from Azure Media Services origin servers by Akamai CDN and distributed throughout its network to viewers playing back NBC Olympics content on Windows, MacOS, iOS, Android and Windows Phone devices. Importantly, Azure Media Services automated the task of spinning up or spinning down origin servers for each stream as necessary to handle rises or falls in concurrent viewers. Akamai picked up feeds from both the US East and US West Azure data centers to deliver them optimally for playback.

Once a live stream was stopped, Azure Media Services automatically applied pre-defined filters to adjust the live programming to video on-demand format, and the content immediately became available for viewing as a VOD program.

Protection and Playback

Content protection is an important aspect for any video workflow in order to ensure that the content provider is complying with applicable business rules, and to prevent piracy and maximize revenue. Aspects of protection include access control, authentication, digital rights management and secure storage and workflow processing. Most professionally run cloud processing solutions ensure that data is secure within their data centers. Content is also typically encrypted in transit during ingest. On the outbound side, content providers have the option of choosing a full-fledged DRM system leveraging typically AES-128 encryption with user-specific rights management, or choosing simpler session-based protection. Content providers will typically leverage more aggressive DRM protection for high-value video on demand content provided under a transactional- or subscription-based model, while typically relying on session-based protection for advertisement-supported content.

In the case of Sochi, NBC Sports chose to use CDN token authentication in addition to Adobe's TV Everywhere authentication to prevent unauthorized access to video streams. For other types of events, such as concerts or pay-per-view subscriptions, content protection solutions such as PlayReady DRM can provide additional layers of protection, such as preventing content recording or restricting media playback to specific secure outputs.

Playback typically falls outside the purview of the cloud-based workflow, which hands content off to the CDN service and then simply ensures that the origin servers are flawlessly delivering content as requested by end users in the appropriate formats and resolution. Content will nearly always have to be delivered to a variety of

players, and a key benefit a good cloud service will provide is out-of-the-box support for every necessary player and device type. For Sochi, NBC Sports chose Adobe Primetime and Flash for the client experience. Android devices were served by a dedicated app from the Google Play store.

Orchestration, KPI and Telemetry

Tying all the above individual components to each other and to the underlying content management system is a non-trivial task. The quality and resiliency of this orchestration will ultimately determine the level of scalability and reliability that is achieved, while the flexibility of the orchestration is critical to ensuring that a workflow can be optimally customized for the given use case. In the case of Sochi, CMS integration was done by Aventus and included ad management, file management, and automated start and stop of programs. Aventus, in turn, interfaced with the orchestration layer within Azure Media Services, which, for example, detected when a program was stopped to automatically switch a live stream to on-demand configuration.

It is critical for cloud-based workflows to have back up at every level, because it's almost impossible to be able to fix something in real time for a live event.

For KPIs, both Aventus and Azure Media Services relied on their own telemetry systems to monitor system performance. Real-time monitoring APIs provided a convenient way to access snapshots of system performance as the events progressed. The teams also relied heavily on numerous dashboards and monitoring logs; for example, an inbound feed dashboard monitored incoming feeds, an encoding dashboard monitored quality of streams as they came out of Aventus and passed into the Azure Media Services packaging platform, while a third dashboard looked into stream health through the storage and origin server stages. Dashboards provided easy access to a diagnostics mode when issues were discovered. Azure Media Services also ensured that their networking partners were available for trouble-shooting any connectivity issues 24x7, given that the winter Olympics were happening half a world away in a significantly offset time zone.

Due to the high-profile nature of the event, additional human monitoring was provided by iStreamPlanet and Microsoft operations teams.

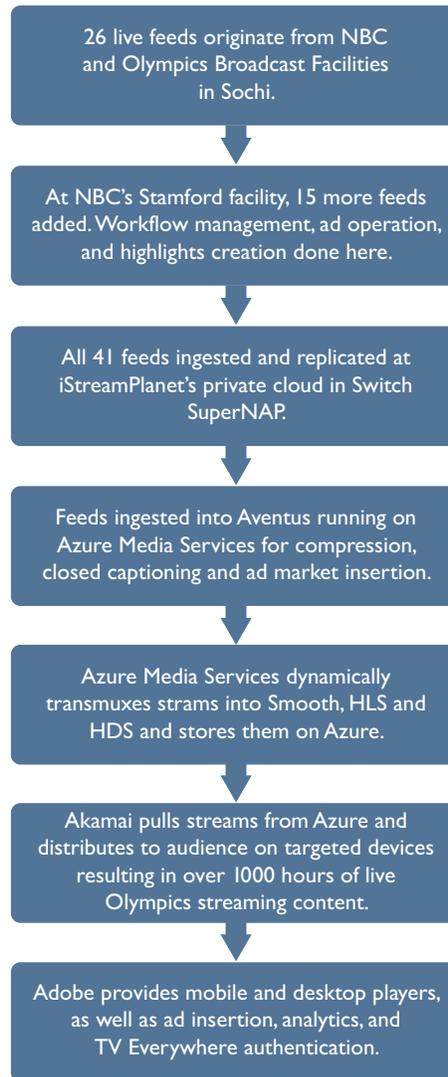
SOLUTION BRIEF: MICROSOFT AZURE MEDIA SERVICES

This paper has so far shown how the cloud has matured into becoming a truly scalable, secure and resilient platform. As we have been using the Sochi Olympics as an example to illustrate the advantages of the cloud, we shall now take a deep dive into the cloud infrastructure used for this use case: Microsoft Azure.

Microsoft Azure Media Services is a cloud-based media processing platform that supports on-demand, live linear and live event workflows. It combines the power and geographical reach of the Microsoft Azure platform with specific component services, including ingest, encoding, protection and serving that are derived from long-standing mature Microsoft technologies such as IIS Media Services, PlayReady DRM and Smooth Streaming.

Furthermore, Azure Media Services is a forward-looking solution that combines the benefits of turnkey cloud-based solutions typically offered as a SaaS (Software as a Service) with the flexibility to craft customized workflows that is traditionally only available by implementing specific workflows with an IaaS (Infrastructure as a Service) provider. Its open architecture and rich set of APIs allow customers to pick and choose

specific components in best-of-breed fashion. The APIs further allow customers to conveniently stitch these components together in a customized fashion to quickly deploy a scalable, resilient, cloud-based workflow that is optimally suited for their particular use case. The breadth of partnerships that Azure Media Services supports is depicted in the Sochi Olympics Live Workflow graphic below.



For porting any production-grade workflow to the cloud, the ability to closely customize the cloud-based workflow and to integrate its configuration and monitoring components with the central production facility is critical to successful deployment. Azure Media Services' compelling value proposition and industrial-grade scalability, resilience and flexibility led it to be chosen as the cloud provider for five independent broadcasters during the Sochi Olympics. Cumulatively across all broadcasters, the platform delivered more than 200 streaming channels across 22 countries with 500 billion storage transactions and 40 million unique viewers – a whopping success story by any standards, and an example that proves that cloud-based platforms have come of age and can reliably power today's mainstream, money-generating online video workflows.

AZURE MEDIA SERVICES: SOCHI OLYMPICS BY THE NUMBERS	
Broadcasters	5 (NBC, CBC, Viasat, América Móvil, NHK)
Data Centers	6
Countries	22 (population of 900 million)
Channels	204
Program Hours	10,000
Channel Hours	88,000
Azure Cores	~10,000
Single Largest Sporting Event	2.0M concurrent
Single Largest Authenticated Sporting Event	2.1M total viewers (NBC numbers only)
Storage	>100 TB
Storage Transactions	~500B
Hours Viewed	~19.3M
Unique Daily Viewers	>100M
Bytes Served	35 PetaBytes

The strongest argument to be made is ultimately in terms of improved monetization. NBC acquired the content rights for the Sochi Olympics for \$750 million. It realized \$800 million in traditional ad revenues, but a significant \$50 million in digital ad revenues. Thus, its online content offerings, delivered in a scalable and reliable way with a pleasing end-user experience, nearly doubled NBC's net revenue from this event.

THE BOTTOM LINE

The growing importance of TV Everywhere has compelled media companies to find technology that will help them get content to their customers where, how and when they want it. This is much easier said than done as the end-to-end components for such a workflow are many, and orchestrating them seamlessly is an incredibly challenging undertaking. Cloud-based models are emerging as the infrastructure approach of choice for online video delivery, as exemplified by the use of the Azure cloud at the Sochi Olympics.

As the cloud continues to disrupt traditional workflows and empowers media companies to fully leverage and exploit its capabilities, concerns about scale, resilience, agility and security are quickly fading. The past few years have seen video consumption being transformed into a nonlinear experience, though the enabling media processes have still remained linear and monolithic. The need of the hour is for media companies to be empowered with the flexibility to deploy their services for live and on-demand consumption across devices and geographies. To survive and thrive in such a competitive environment, the cloud has emerged as the critical foundational element that will catalyze the transformation of the media industry.

Auckland
Bahrain
Bangkok
Beijing
Bengaluru
Buenos Aires
Cape Town
Chennai
Colombo
Delhi/NCR
Detroit

Dubai
Frankfurt
Iskander Malaysia/Johor Bahru
Istanbul
Jakarta
Kolkata
Kuala Lumpur
London
Manhattan
Miami
Milan

Mumbai
Moscow
Oxford
Paris
Pune
Rockville Centre
San Antonio
São Paulo
Sarasota
Seoul
Shanghai

Shenzhen
Silicon Valley
Singapore
Sophia Antipolis
Sydney
Taipei
Tel Aviv
Tokyo
Toronto
Warsaw
Washington, DC

Silicon Valley

331 E. Evelyn Ave., Suite 100
Mountain View, CA 94041
Tel 650.475.4500
Fax 650.475.1570

San Antonio

7550 West Interstate 10,
Suite 400
San Antonio, TX 78229
Tel 210.348.1000
Fax 210.348.1003

London

4 Grosvenor Gardens
London SW1W 0DH
Tel +44 (0)20 7343 8383
Fax +44 (0)20 7730 3343

877.GoFrost
myfrost@frost.com
www.frost.com

Frost & Sullivan, the Growth Partnership Company, works in collaboration with clients to leverage visionary innovation that addresses the global challenges and related growth opportunities that will make or break today's market participants. For more than 50 years, we have been developing growth strategies for the Global 1000, emerging businesses, the public sector and the investment community. Is your organization prepared for the next profound wave of industry convergence, disruptive technologies, increasing competitive intensity, Mega Trends, breakthrough best practices, changing customer dynamics and emerging economies?

For information regarding permission, write:

Frost & Sullivan
331 E. Evelyn Ave., Suite 100
Mountain View, CA 94041