# The Kollective SD ECDN

HOW IT WORKS

## The Kollective Software Defined Enterprise Content Delivery Network (SD ECDN)

A *software-based network* that orchestrates both an enterprise's network infrastructure and its end-user devices into an adaptive, continuously optimizing, fully distributed content cache and delivery system. Its formation and operation are fully software-defined, providing the flexibility, agility, and central control commonly afforded by software-defined systems.

## The Corporate Network Challenge

The typical deployment context for the Kollective SD ECDN is a large, multi-national corporation with a globally distributed workforce, depending on a substantial but heterogeneous corporate network. The diagram below is representative of this concept: a high-capacity corporate backbone in the home country with lower capacity in-country backbones and links to branch offices, fanning out to sometimes often very low-bandwidth WAN links in remote offices.

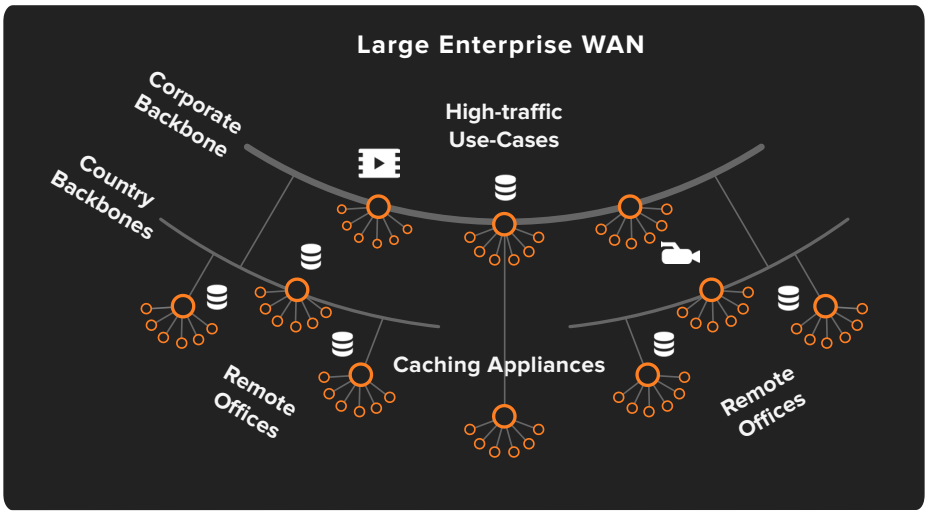

As more and more business functions become IP-based, demands on a corporate network's capacity increase to the point where the network is a constraining and contested resource. Some scenarios that are particularly problematic include: the release of a new training video on an internal portal that will be in high demand in the branch offices, or an important all-hands webcast from the CEO. These generate substantial "north-south" traffic, from the backbones out to the edge, which can easily result in saturated WAN links and the disruption of critical business functions. These practices are often either banned outright, or require the purchase and deployment of many expensive hardware caches, WAN optimizers, streaming-server repeaters and other devices to reduce this north-south traffic over congested WAN links.

## The Kollective Difference: Unique and Efficient Software Delivery

The Kollective SD ECDN addresses these content-delivery challenges entirely with software, leveraging existing network infrastructure, as well as latent but generally unused capacity in the broader infrastructure, notably storage and serving bandwidth on end-user devices, to easily handle these cases. The Kollective SD ECDN is a set of Kollective-managed, cloud-hosted control and origin servers and a small software agent deployed on employee devices throughout the company, as shown below.

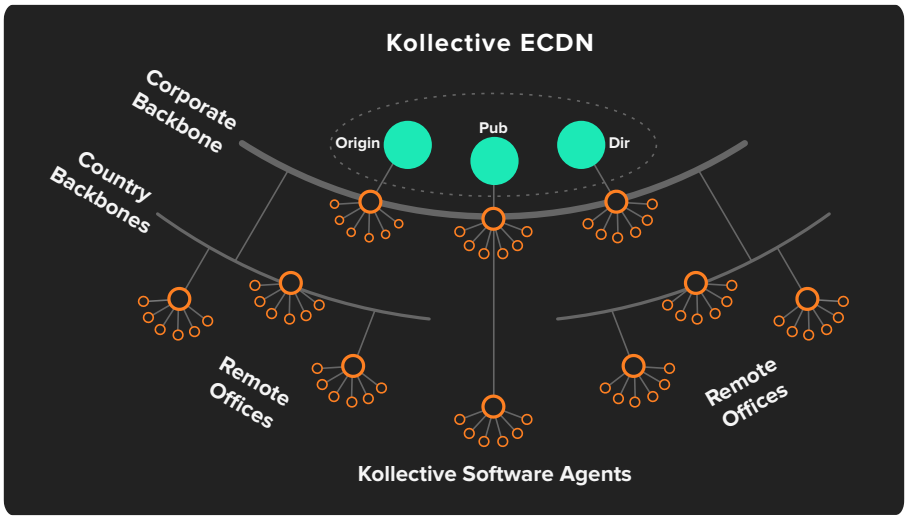


The central servers and the agents *collectively* form an adaptive and distributed content delivery and caching system to ensure that upwards of 99% of content is delivered via controlled, localized, traffic that doesn't congest WAN links. All of these software components cooperate to deliver content, secure it via a multi-layered crypto framework, and form an optimal delivery overlay mesh that dynamically adapts to network changes. All aspects of its operation are software-defined.

## Kollective SD ECDN Architecture and Key Components

The Kollective SD ECDN architecture exhibits a classical three-layer SD structure:

1. An **Application Layer** comprising management, analytics and content applications, all built on a set of north-bound APIs provided by the second layer.
2. A **Control Layer** that centrally orchestrates and manages the network. It provides the high-level API for the application layer and uses a common set of south-bound protocols to manage the third layer.
3. A **Data Layer** comprising all the components that are used to form the Kollective SD ECDN's network - a small number of central, cloud-based Kollective head-end servers, the existing corporate network itself, plus the end-user devices running the Kollective agent.

The top two layers, Application and Control, and portions of the Data layer, are hosted in the Kollective Cloud providing a fully-managed SaaS solution, only the Kollective agent needs to be deployed on end-user devices within the enterprise.

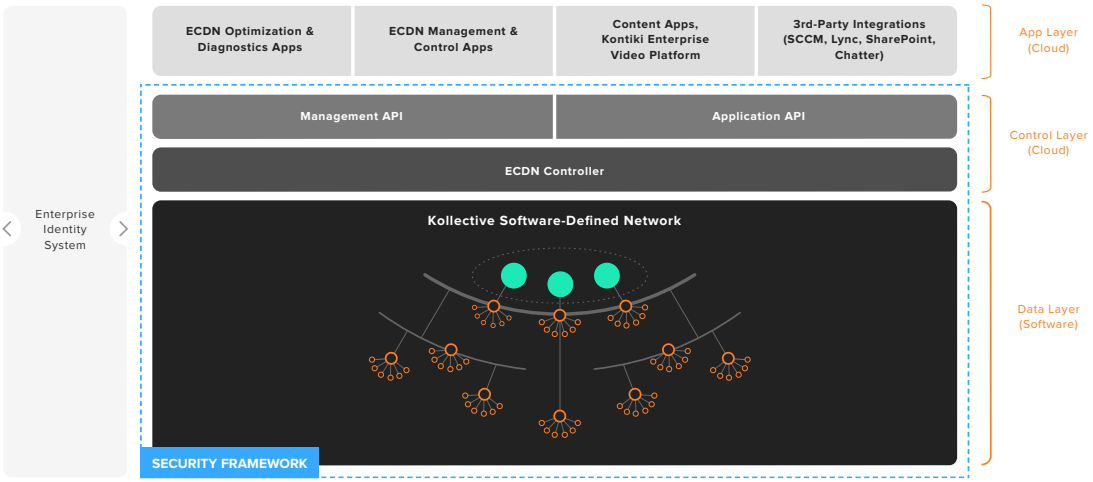


The **Application Layer** contains:

- **Management and control applications** used to manage and monitor SD ECDN operation and orchestrate per-enterprise network policies and operational rules. There are over 250 software controlled parameters covering these policies and rules, ranging from protocol and port preferences, through hierarchical location and device group definitions, in order to provide full control of network formation and traffic patterns, as well as bandwidth, disk and CPU use caps for individual end-user devices.
- **Delivery monitoring and readiness-testing applications** providing exhaustive pre-flight and post-flight analytics. The network readiness testing application is a unique aspect of the Kollective SD ECDN, taking advantage of the deployed agents to perform automatic and invisible delivery testing. This provides crucial feedback and operational confidence prior to an important delivery event, such as a live CEO webcast.

- **Content-specific applications** that provide user-focused capabilities for various content-delivery use cases and leverage the unique delivery capabilities of the Kollective SD ECDN. The available applications include a suite of video-centric solutions offered by Kollective under the Kontiki brand:
    - **Kontiki MediaCenter** – a feature-rich, fully customizable video portal platform to build enterprise intranet "YouTubes" for business applications.
    - **Kontiki Webcaster** – an easy-to-use webcasting platform that can be set up and run by individual departments, utilizing the Kollective SD ECDN's live-streaming capabilities to broadcast webcasts to thousands of concurrent viewers across the enterprise.

  **Kollective SD ECDN** is integrated with several partner applications, including:
    - Microsoft Skype for Business
    - InterCall IWS
    - Nasdaq Webcasts
    - Panopto VCMS

The **Control Layer** contains:

- **Network directory servers** that keep track of the delivery network topology and content disposition within the network. They provide key delivery-mesh formation intelligence to the network as a whole.
- **Agent Directors** that monitor and manage the agents, pushing any agent-specific software-defined network controls out to the agents, managing automated subscription and targeted deliveries and assisting the network readiness-test system in orchestrating agent test sets and runs.
- **A Content Management System** that provides fully authenticated content ingestion, transcoding, encryption, metadata and access controls. This system integrates with an enterprise's authentication servers to provide enterprise directory-compliant content access controls.
- **Live video stream ingest endpoints** that can operate in either push or pull mode using either external encoders or other live video streams, conditioning and directing the streams for delivery out through the Kollective SD ECDN.
- **Network readiness-test managers** that are used to set up and monitor pre-flight readiness tests. These allow various kinds of delivery events to be tested using either explicitly selected sets or statistically sampled sets of end-user devices.
- **SD ECDN status and monitoring servers** that take in constant status and monitoring data from all the components, enabling monitoring and analytics services. These take in delivery event and network performance data from each agent, providing both for a global view of the network's operations, as well as content usage analytics for content creators.

The **Data Layer** contains:

- **Delivery network origin servers**, hosted in the Kollective Cloud. They contain the source copies of any on-demand content and originate all live streams delivered through the SD ECDN. Kollective's delivery-mesh formation algorithms work to minimize traffic from the origins, which usually act as single-copy originating sources for delivery meshes. These meshes are formed from agents within the corporate network itself, or as guaranteed copies of last resort if needed.
- **Kollective agents** running silently in the background on enterprise end-user devices and desktops. They cooperate with control layer components, the origin servers, and with one another to form the adaptive, distributed delivery network and edge cache.

## Kollective SD ECDN Benefits

- **Economical** – No additional hardware needs to be purchased, deployed, managed or upgraded; resulting in capital and operational cost savings.
- **Minimal Deployment** – Only the small agent needs to be deployed within the enterprise, typically via a desktop-management system. This is much simpler than deploying distributed hardware solutions and can often be accomplished in a matter of days.
- **Adaptive** – Automatically and dynamically adjusts to changes in traffic patterns and physical changes in the underlying network.
    - **Self-scaling** – The more requestors for content there are, the more resources are available for distributing the load.
    - **Self-healing** – If a node stops serving, others take over automatically as necessary.
- **Intelligent** – Enables capabilities such as background push delivery and live-event readiness testing, as well as future network edge monitoring & control applications.
- **Centrally Controlled** – Being an SD ECDN, all operational aspects are managed through a SD ECDN Controller.

## The Kollective SD ECDN in Operation

### Trust establishment

Once the agents are deployed and activated, they perform a lightweight discovery process, first contacting the central SD ECDN control servers to establish a trust framework based on 1032-bit X.509 certificates. Every node in the network, along with each central server and end-user device, is allocated a unique certificate containing a PKI key-pair, the public key of which is used as the node's main identifier within the network. All messages sent between nodes are signed by the sender and encrypted for the receiver using these certification keys. The central server nodes' certification keys are signed by the Kollective certificate-authority, thus assigning to them authoritative, system-server status. By leveraging this trust framework, no malicious commands or content can be introduced.

### Topology Discovery

The agents then perform a configurable sequence of topology discovery probes. This process includes a traceroute to the central servers and gateway router, local NIC inspection, and LAN or subnet broadcasts or multicasts. The results are sent to the ECDN Controller to build a global topology graph. The agent also keeps this information locally so that it knows its own neighborhood. This discovery process is repeated whenever a device restarts so that any changes can be noted. In addition to this startup process, each node, sends a periodic status report to the central servers that contain the latest topology discoveries, available content listings, and various delivery and network metrics, all of which help with the formation of optimal delivery paths during actual content delivery.

### Content Publishing

The Kollective SD ECDN is a fully-managed ECDN, meaning all content publishing and live event scheduling is authenticated and secure. A user authorized to publish connects to the SD ECDN through one of a number of content-management portals or APIs and can then perform several tasks, such as:

- Creating a logical content item that can be associated with one or more physical files or streaming sources, (typically as alternative formats, sizes, or bitrates) so that the consuming agent can pick the best format for its local context.
- Adding descriptive metadata or portal-specific structure such as text descriptions, thumbnails, keyframes, channel location, and more.
- Defining content subscriptions and feeds that enable automatic background downloads.
- Defining availability date ranges or live event schedules.
- Setting up end-user access controls; Kollective has a sophisticated content security system that integrates with the enterprise's own identity services.

Content items are assigned a unique location-independent identifier that can be embedded in the Kollective SD ECDN URLs. These are commonly made available to users as clickable items in a content portal.

Content data itself is ingested into the Kollective SD ECDN in a number of ways, including HTTPS upload for static files or push/pull stream endpoints for live streams. Depending on ingest mode and publisher instructions, the content data may be transcoded, virus-scanned and encrypted, and in all cases has a set of data-block cryptographic digests created that will be used later during delivery to validate content as it arrives at a receiver.

## End-user Authentication and Content Requests

In most cases, content stored in the Kollective SD ECDN is access-controlled and requires end-user authentication to make it available for delivery. The SD ECDN provides a number of authentication modes, including simple username and password and several single sign-on schemes that can interface to an enterprise's authentication system over protocols such as LDAP and SAML. Once authenticated, the SD ECDN generates a time-limited token for the user that securely encodes the user's credentials and group membership.

Content can be requested explicitly, by presenting a content URL to the agent's localhost HTTP or RTMP server as a clickable link in a content portal, or implicitly if the user has subscribed to content feeds or subscriptions. In the latter case, the agent manages subscriptions automatically in the background, downloading content under the control of the subscription publisher's policies. This makes content available either in a local directory or via the agent's localhost server using the content item's localhost URL.

In both cases, the agent presents the user's authentication token and the content identifier to a system server that checks access rights and returns encrypted content metadata, block digests, and a secret download ticket. This download ticket is used to securely request content fragments from other nodes in the network during delivery and the block digests are used to validate the fragments as they arrive.

## Delivery-mesh Formation and Content Delivery

The Kollective network uses a proprietary protocol, known as Kollective Delivery Protocol (KDP), which is specifically designed for distributed delivery and built from the ground-up on a PKI security model. It can be carried over UDP, TCP or HTTP, and will automatically choose the best carrier for a given context. The UDP-based version is particularly efficient and supports software-defined quality-of-service settings using its adjustable congestion-control capabilities. The key benefits of using KDP are:

- TCP-like reliability.
- Enhanced congestion avoidance and dynamic throttling so the agent can throttle back and allow business-critical traffic to flow uninterrupted.
- Knowledge of the live stream's minimum throughput to sustain a good viewing experience, regardless of latency.
- QoS controllability.

In general, an agent requesting delivery will attempt to get different fragments of a content item in parallel from as many nodes in the network as it can find, subject to software-defined topology boundaries (connection limits and bandwidth caps) and then bond the bandwidths of the available servers to speed up delivery.

To find available source nodes, the agent begins a source-discovery process that is repeated during delivery to adapt dynamically to network and resource changes. As more nodes request the same content, a delivery mesh emerges, with nodes collectively pipelining, caching and serving various parts of the content for one another. This process adaptively seeks an optimal mesh, maximizing local, east-west traffic and effective serving bandwidth, while minimizing north-south, WAN link traffic and in this process lies the essential value of the Kollective SD ECDN.

The mesh formation is a cooperative process between the central directory servers and agents. The directory servers have a dynamic, global view of network topology and content disposition, based on content request history and the periodic readiness-reports sent by the agents. The requesting nodes continuously discover and evaluate sources by:

- Getting a list of candidate sources from directory servers prioritized by proximity and other metrics
- Sending and receiving local content-discovery broadcasts or multicasts
- Receiving content requests from other nodes

Each node rotates through its prioritized sources, making multiple concurrent connections, discarding poor sources, and re-engaging source discovery as needed, all under the control of software-defined formation policies, such as LAN-focusing, topology boundary rules, throttling rules, and so on.

Serving requests are only honored if the requestor supplies a valid delivery ticket, which is obscured by hashing it against the requesting node's ID to prevent ticket hijacking. Received content is accepted only if it passes block-digest tests. Content data blocks sent between nodes are encrypted using unique, ephemeral 128-bit symmetric keys that node pairs establish on connection.

The Kollective SD ECDN supports a range of delivery modes and tunes policies for mesh-formation, block requests and traffic-control to best suit each mode.

## Background File Download

Requesting nodes choose random blocks to download so they can cross-serve one-another to reduce load on origin servers and WAN links. In addition to using standard QoS/DiffServ controls, the KDP software measures and benchmarks the packet round-trip and dynamically throttles the download speed. It makes background download traffic deferential to all other network traffic, effectively making the download soak up idle bandwidth. Kollective SD ECDN agents detect user activity and politely throttle CPU and bandwidth use, so as not to interfere with foreground tasks on the device.

## Video on Demand Streaming

During video stream playback, blocks in the buffering region ahead of the playhead are requested to ensure smooth playback, falling back to random block requests if the buffer is well filled. QoS levels are typically set to compete fairly with other traffic.

## Live Event Streaming

All viewing nodes effectively want the same portions of the stream at the same time and so nodes within a locality cooperate to elect a well-performing lead node that will get a single copy of the stream across the WAN link and then pipeline it out though a mesh formed from the other local nodes. The leader-node election itself is adaptive and leadership can be handed off to better performing nodes dynamically during an event. QoS is set high for important live events to ensure smooth event viewing. There are two additional components that contribute to increasing the QoS:

- Knowledge and enforcement of the stream's 'minimum' bitrate establishes a latency immunity.
- The Adaptive Bit Rate mechanism finds the optimal bitrate for each WAN connection, resulting in the best possible viewing experience.

## Reporting and Analytics

All nodes make periodic reports to a central analytics system, containing delivery event details, local loading and serving metrics, video playback stats and data about other delivery-related activity. This allows the Kollective SD ECDN analytics reports to be produced both on content delivery and use, as well as network efficiencies and performance.

## Integrating the Kollective SD ECDN: Kollective for Skype Meeting Broadcast

Kollective SD ECDN integrates with numerous enterprise software partners and is highly extensible through a robust API built to industry standards for speed and security. An excellent example of this integration is Kollective for Skype Meeting Broadcast. Kollective for Skype Meeting Broadcast solves enterprise video delivery challenges by routing secure, high-quality live video on top of your existing network infrastructure. This integration showcases how the Kollective SD ECDN is particularly adept at solving network congestion problems.

The Kollective SD ECDN is accustomed to solving internal congestion issues, primarily at WAN links. When all the video traffic is coming in over the internet gateways, as in the case of serving streams from the Azure CDN, the enterprise internet gateways also become a potential bottleneck in addition to the internal the WAN links.

By default, Skype Meeting Broadcast generates video sessions through the Azure Content Delivery Network (CDN) and delivers the video experience through the Azure Media Player running on a web browser. In order to scale meetings up to 10,000 users, IT support teams provision the Kollective SD ECDN integration. The availability of the SD ECDN delivery path is transparent to individual users, and the Kollective SD ECDN is activated as needed. The high-level architecture of the Kollective SD ECDN integration with Skype Meeting Broadcast is described below:

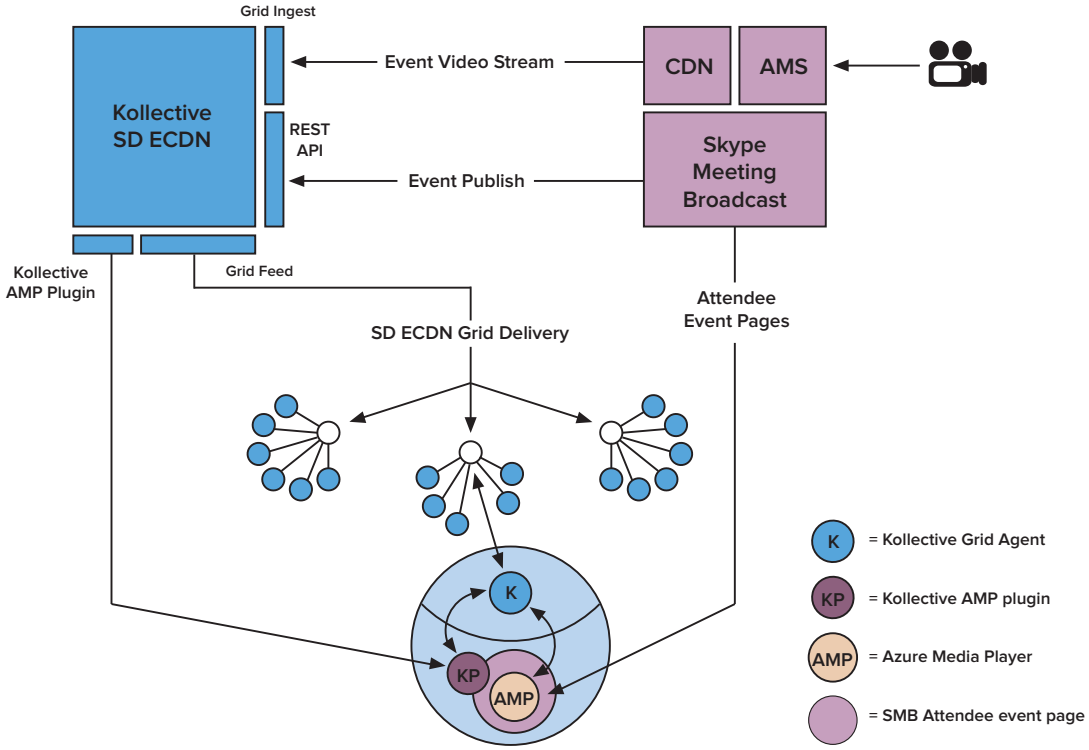


**Skype Meeting Broadcast / Kollective SD ECDN Integration**

:

## The Kollective SD ECDN:

- Is fully compliant with Azure Media Services streaming requirements supporting all streaming protocols – SmoothStream, MPEG-Dash and HLS
- Supports pass-through of AES encrypted streams for highly-secure enterprise applications
- Is integrated through Kollective's secure REST API on the server side, requiring just a single call to publish content and event streams into SD ECDN
- Is integrated via the Azure Media Player SDN plugin framework on the agent side
- Provides completely transparent delivery of Skype Meeting Broadcast event streams

An example process of setting up and running a Broadcast in Skype Meeting Broadcast:

1. Upon live event creation and activation, Skype Meeting Broadcast sends Kollective SD ECDN information about the event securely, including metadata and streaming information.
2. Attendee navigates to a Skype Meeting Broadcast event page. The user device contains:
    - Skype Meeting Broadcast app with embedded Azure Media Player and reference to Kollective's Azure Media Player plugin
    - Kollective SD ECDN agent
3. Attendee requests playback
    - Azure Media Player invokes Kollective SD ECDN plugin which checks if an SD ECDN agent has been detected.
    - If detected, prepare Azure Media Player for streaming from the SD ECDN agent.
    - If an agent is was not detected, the player uses the Azure CDN.
4. Kollective grid pulls the stream from the Azure CDN and deliver it throughout the smart grid.
5. Kollective SD ECDN plugin for Azure Media Player reports play statistics back to Kollective.

## Summary and Key Benefits

The Kollective SD ECDN offers an end to end delivery solution for video, software and other large files. It's built to be used today but in the future the type of content and supported applications will continue to expand.

### Benefits Summary:

**Kollective Solves some of the Biggest Network Challenges in the Enterprise**

- Stream a high quality, live video All Hands meeting to all employees reliably, without impacting the network.
- Video-enable enterprise applications, like SharePoint or the corporate intranet with thousands of videos managed centrally in one platform.
- Move large files around your network with ease. Need to make a 4GB Microsoft Office Update available to all employees in India? No problem.

**Kollective Surpasses Enterprise Expectations by Utilizing Breakthrough Technology**

- Software-Defined Technology - Kollective's SD ECDN acts as an intelligent network. Every computer is a content server.
- Control Layer – Network becomes highly configurable: characteristics of the network functions are configured via software to determine the key attributes of the network's function.
- Adaptive Response – Guaranteed most efficient, timely, and complete delivery; dynamically redistributes load based on network changes within the guidelines set by the SD ECDN Controller.
- A highly extensible and robust API enables an integration with Microsoft that scales Skype Meeting Broadcast to 10,000 employees simultaneously.