Summary Report
On
Akamai Streaming

Allan Lee

Akamai streaming is a service offered to content providers to enhance the streaming of their content to users (customers) in such a way that the end product is an accessible, steady, reliable, and hence satisfying experience for the end user. The Akamai streaming service came out of a solution to some of the drawbacks with the internet's infrastructure as it became increasingly vulnerable with the rapid growth of websites and the number of users trying to access those sites. From a content providers perspective, websites can become overwhelmed from request loads, which overstress some of the site key infrastructures such as front-end web servers, network equipment, bandwidth, or the back-end transaction-processing module. The Akamai streaming service came out of research efforts aimed at solving the above-mentioned website vulnerabilities, referred to as the "flash crowd problem".

The Akamai approach to the solution of the "flash crowd problem" is based on the observation that web content served from a single location can suffer from serious problems which impacts site scalability, reliability and performance. Based on this observation, a distributed system was devised to serve site requests from a variable number of surrogate origin servers located at the network's edge. It is thought that by caching content at the Internet's edge, the demand on sites infrastructure can be reduced and since edge servers will be closer to the user, faster service can be provided. As of October 2002, Akamai system has grown to over 12,000 servers in more than 1,000 networks. Operating these servers in the various locations poses many technical challenges, including how to direct user requests to appropriate servers (mapping), how to handle failures, how to monitor and control the servers, and how to software across the system. This summary explores the technologies and the architectures involved in the implementation of Akamai streaming service and the approach to these challenges.

Akamai has a large deployment of servers that forms a fault-tolerant network for streaming media for their customers (content providers). These servers are collocated within data centers and points of presence (POP) of major Internet backbone carriers and Internet service providers including satellite providers and broadband networks such as cable and DSL. Each deployment of servers is referred to as an Akamai region. The distributed network is designed with a capacity that intentionally and substantially exceeds the expected demand. The network can provide hundreds of thousands of streams simultaneously. Akamai uses a propriety technology called SteadyStream^{SM} to deliver live and on-demand streaming content from the server at the entry point to its many servers located at the edge of the Internet.

For delivery of broadcast streaming, once the stream is in the Akamai network, the challenge is to transmit the content from its single point of origin to the networks edge with minimal loss or latency. To

achieve this, at the Akamai entry point there are two servers for redundancy, running reflector software, each being capable of handling many streams from multiple content providers. Once in the network, copies of the streams are then rebroadcast to predefined reflectors via UDP. For example, Windows Media Technologies (WMT) encapsulated in RTSP encapsulated in UDP over IP. The reflectors are so positioned as to ensure, from a network and geographical standpoint, a highly fault-tolerant network. Each reflector in turn, rebroadcasts its copy of the stream to its subscribing region. Each subscribing region then receives multiple copies of the same stream, with each stream taking a different route over the Internet. The idea being that congestion on different routes, resulting in lost packets, is unlikely to impact each stream equally. In each Akamai region, there is a mechanism that recreates the original version of the stream in real time from the multiple received streams. The algorithm execution takes the first packet #1 that it receives and discards the rest, takes the first packet #2 and discards the rest and so on. The reassembly mechanism within a region then makes the stream available to every Akamai streaming media server in that region. When a user clicks on an Akamaized stream, the stream delivery takes place from the nearest available edge server.

For on-demand streaming, the Akamai network employs a pull architecture in which content is replicated in streaming server caches in response to user requests. The Akamai ARL uses serial numbers to enable the network to load balance among the streaming servers to ensure that no single server is responsible for the delivery of the entire content. Serial numbers also allow servers to respond in real time to customer requests and adjust dynamically to meet the end-users varying demands. SteadyStream^{SM} uses HTTP protocol to transfer on-demand streams to its edge servers. This is contrast to the use of UDP for live streaming.

Akamai streaming supports the three major streaming formats which includes, Apple Quick Time (QT), Microsoft Windows Media Technologies (WMT) and RealNetworks RealSystem G2 (Real) for both on-demand and broadcast streaming. To use the service, the content provider has to ‘Akamaize’ the URL that point to the file containing content that is to be streamed. Basically, Akamaizing pre-pends the streaming URL with a string containing an Akamai domain and some additional control information, resulting in an Akamai Resource Locator (ARL). Two examples ARLs are shown below for live and on-demand audio or video for the URL rtsp://www.foo.com/movie.mov. These are shown for WMT and Real formats. The corresponding ARL looks like this:

On-demand:

WMT: mms://a9.m.akastream.net/7/9/801/v001/www.foo.com/movie.asf

Live or Broadcast:

Real: rtsp://a9.r.akareal.net/live/D/9/801/v001/reflector:35001

The ARL contains metadata that instructs the Akamai server how to process the request. The processing is facilitated by Akamai’s EdgeAdvantage, a control software technology that resides in each of the DNS

located throughout the network. Information such as, load balancing metrics, special signal processing rules, identifying the customer for billing purposes, and when to refresh the content in the cache as well as where to get the content in case the optimal edge nodes do not have it, are some of the control information contained in the metadata. This control data facilitates intelligent mapping, which resolves a hostname based on several factors, including, service requested, user location, server load, server health, content requested and network status.

Load balancing of the system is achieved through continuously monitoring the state of services, and their servers and networks. Akamai monitors the system health through the use of agents that simulate end-user behavior by downloading web objects and measuring their failure rates and downloading times. This information is used to monitor overall system performance and to automatically identify and bypass problematic data centers and servers. This monitoring capability makes the entire system highly resilient to catastrophic network events, since it can respond to these events in seconds.