

Networked Systems Research at Akamai

Erik Nygren
Akamai Technologies
nygren@akamai.com

Ramesh K. Sitaraman
Akamai Technologies
& UMass, Amherst
ramesh@cs.umass.edu

Joel Wein
Akamai Technologies
jwein@akamai.com

ABSTRACT

In this special section of OSR, we present a selection of research papers that relate to Akamai's platform and technology.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Distributed networks; C.2.4 [Distributed Systems]: Distributed applications

Keywords

Distributed Systems, Network Architecture, Content Delivery

1. INTRODUCTION

Akamai was founded more than a decade ago with the vision of helping enterprises effectively deliver rich content and applications on the Internet. Akamai's distributed platform currently consists of more than 61,000 servers deployed in over 70 countries and 1000 networks. The platform delivers roughly 15-20% of the web content world-wide and has powered some of the largest events in the history of the Internet.

Since Akamai's founding, there has been a rapid evolution in its technology and services. More than a decade ago, Akamai pioneered the concept of a Content Delivery Network (or, CDN) to accelerate web content. Since those early days, a series of innovations has quickly expanded the capabilities of its platform into radically new domains. Today, the platform supports application delivery that can accelerate entire web or IP-based applications, media delivery that provides HD-quality delivery of live and on-demand media, analytics technologies that give our customers visibility into how their content is used, and EdgeComputingTM technology that deploys and executes entire Java J2EE applications in a distributed fashion.

Since its academic origins at MIT, Akamai's technology and product innovation has been closely-tied to advances in basic research in computer science. The close interaction between research and development at Akamai is truly symbiotic. Several of the unique challenges in building Akamai-scale services on the Internet have been the motivation for fundamental research published in CS conferences and journals. In turn, many of the recent research advances in distributed systems and networking have influenced the way Akamai builds its products and services. The close interplay between research and development is fostered partly due to the fact

that a number of Akamai engineers and technologists are themselves established researchers. Further, a number of academics who were historically associated with the founding and launch of Akamai have continued to spend a part of their time on Akamai-inspired basic research.

2. IN THIS ISSUE

We start this collection of research papers with an architectural overview of the Akamai platform. The overview starts with the challenges that enterprises face in delivering content and applications on the Internet and sets the context for how the Akamai platform addresses these challenges. Then, the paper overviews the design principles that underlie the platform, which is in essence a distributed system of massive scale and reach. Next, it describes much of the platform architecture, followed by further details on the subsystems that work together to effectively provide both a high-quality experience for end users as well as the control and visibility required by enterprise customers. Finally, the paper includes several case studies of enterprises using the platform as well as references to research literature for further reading. The paper is meant to be a one-stop reference for the state of the art that we hope seasoned researchers and students alike find useful.

We follow this paper with an overview of Query, an important subsystem of the Akamai platform. Query provides a near-real-time relational database system for network monitoring, and has seen a decade of continual use and evolution. With the increased emphasis for providing real-time visibility into Internet events, a description of this subsystem is timely.

Due to Akamai's widely distributed footprint, it is in an excellent position to collect valuable data about the Internet itself. The richness of this data set has been a key driver for a number of research initiatives and publications over the years. Selected excerpts of this data are presented in near-real-time on the Akamai website (<http://www.akamai.com/>), as well as in the quarterly State of the Internet Report (<http://www.akamai.com/stateoftheinternet/>). Our third and final paper presents insights into how the Internet is evolving based on data gathered from across Akamai's global server network, including trends in attack traffic, connection speeds, Internet penetration, and broadband adoption.