ORACLE COMMUNICATIONS

Advances in Communications Applications Servers

Driving New Innovation into Communications

ORACLE WHITE PAPER | JANUARY 2016

ORACLE

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Executive Summary

The proliferation of IP-enabled mobile and fixed devices capable of offering Web, telecom and entertainment services, such as Voice over LTE(VoLTE) smart phones, TVs, home appliances, gaming consoles, and automobiles are causing operational and business challenges for Communications Service Providers (CSPs) and Enterprises. CSPs are challenged to reliably deliver innovative IP-based communication services with maximum profitability. At the same time enterprises are looking for services to achieve new levels of productivity. New web capabilities driven by Web Real Time Communications (WebRTC) are opening the door to a new wave of advanced real-time communication applications. Additionally, Network Function Virtualization (NFV) promises to bring the agility and scale to the increasing demand of communications services, which has a significant impact on how services are created and deployed. Deployed communications services in the network is often proprietary service logic often in the form of computer programs that run in software application servers distributed across several hardware platforms.

In order to overcome these challenges, CSPs and Enterprises are migrating away from expensive, closed, proprietary and application-specific legacy platforms, and moving towards low-cost, open, standards-based Communication Application Servers (CAS), which dramatically lowers the time and cost of adding new features and extensions to existing IP-based communication services. With this approach, services are developed and deployed as one or more computer application on an open standard platform, and those applications are re-used across multiple networks, whether it's the Web-domain, mobile, broadband or fixed networks.

The CAS is a general purpose computing platform processing real time voice, video, and data communications service logic while reliably supporting millions of subscribers or end-points. The benefits and features delivered by Oracle in the latest CAS include:

» Rapid service innovation, integration and deployment: Advances in programming converged communication services, as defined by the Java community, provides agile, highly adaptive protocol interfaces and state event processing.
» Lower cost of ownership: High availability, highly performing, dynamic scalability clustering with data grid provides high speed, highly reliable data storage with built in event processing.
» Agile, scalable Virtualized Network Function (VNF): The virtualized CAS may be orchestrated by a Service or Application Orchestrator as defined by the NFV standards.

The advances in the CAS is due in part to the convergence of IT with Communications, where IT provided an open, standard, cloud environment allowing portable applications to be quickly taken to market in a multi-threaded container-based environment. The communications industry requires reliable, interoperable platforms where sessions and protocols interacted with real time communications on a large scale. The convergence has brought open, reliable, CAS that integrate with all networks allowing portable efficient service creation.

Rapid Service Innovation, Integration and Deployment

The foundation of a CAS is the reliable execution of services based on the state of network driven by protocol events.   Session Initiation Protocol (SIP) is the industry standard for mobile, fixed, and IP communications.   Servlet technology encapsulates the Java event driven model providing rapid development of event driven processes. Converging SIP and Servlet technology provide a powerful programming environment where the average Java developer may create communications services for carrier grade networks, enterprise communications, or internet communications   The latest release of the SIP Servlet standard, version 2.0 as defined by JSR 359, introduces highly adaptive protocol interfaces that allow the mixing and matching of any protocol to control and manage a communications session.   A typical network connection may use SIP to set up and tear down the session, diameter to log or rate the session, and media server interfaces for enhanced voice and video processing.   A CAS application, which may be defined as a “service building block” (SBB) in telco parlance, invoke protocols that have been added or defined in the JEE container.   This programming model provides a modular architecture for building communications services and agile adaptive interfaces to connect and dialogue the service to any network element.

“An application server (AS) is a network element capable of hosting and executing IP multimedia services based on the Session Initiation Protocol (SIP) signaling. … the next-generation SCP was carried forward to embrace IP transport and SIP signaling to support other IP-based multimedia services, notably (but not limited to) video”

DAVID SNOW

PRINCIPAL ANALYST FOR SERVICE PROVIDER INFRASTRUCTURE,

CURRENT ANALYSIS

The technical enhancements to Sip Servlet 2.0 include:

» Highly adaptive agile protocol interfaces.   Defined in JEE, Plain Old Java Objects (POJOs) offer reduced boiler-plate code with well defined feature rich annotations.   The POJO simplifies programming by removing the complicated java rules of inheritance, interfaces, and other java constructs that were hard to debug and made java code less portable.   An Annotated POJO adds a “@” to the POJO.   While SIP Servlet 2.0 is backward compatible with SIP Servet 1.0, supporting programming constructs such as the following example –

```
protected void doInvite(SipServletRequest req) throws …
```

The SIP Servlet 2.0 annotated POJO construct looks like

```
@Invite
        public void incomingCall(SipServletRequest request) {
```

The Annotated POJO is an event driven model similar to servlets, without the rigid definition of a servlet.   The developer has greater flexibility to closely model the business logic and add in other protocols such as diameter for rating or billing, or JSR 309 for media processing

- » *Stateful Processing* – the public switch network has been described as a very large state machine – tracking, maintaining, and building the state of phone calls. In addition to POJOs, JEE defines Context Dependant Interjection (CDI) - a programmatic interface for building and maintaining stateful objects. CDI is a data programming construct that adds stateful processing to data objects. Built into a CDI objects are controls for lifecycle support, timers, events… all the Java controls required to build and maintain a connection in a network state event objects.
- » *Concurrent Session Management* - In communications, feature and service interaction becomes complex when mixing and matching applications to build a service. Services such as Conference calls, web collaboration, or combining subscriber features, etc become complex to manage concurrent sessions. While not eliminating the complexities of service interaction, SIP Servlet 2.0 provides the standard management controls over active concurrent sessions.
- » *Secure Duplex Connections* – the public switch network was built on secure full duplex lines, that is, 2-way voice communications controlled by a CSP. The internet has finally caught up with WebSockets - secure 2-way communications connection. A WebSocket allows for 2-way real time voice, video, and data communications.
- » *Standard open JEE constructs* – there are many other capabilities a CAS may use from an JEE environment, such as communication objects through JSON, XML processing JAX, communications through JMS.

One of the largest cable providers took advantage of SIP Servlets and the latest CAS features by rapidly building highly adaptive agile SIP Servlet building blocks, integrating the SIP Servlet building blocks into a set of customer services, and deploying the services to their customers on top of a IP Multimedia Subsystem(IMS) network. While IMS is designed for mobile phone networks, IMS provided the cable company a well defined infrastructure for deploying new services while maintaining a very large subscriber base. The primary objective for building applications in the CAS in IMS was to run any customer based service on any device – a set top box, a mobile phone, a tablet, a TV - and converge TV viewing with phone calls. Their customers could receive calls while watching TV, send the call to another device or to voice mail, or any number of calling services through the cable set top box or mobile device.

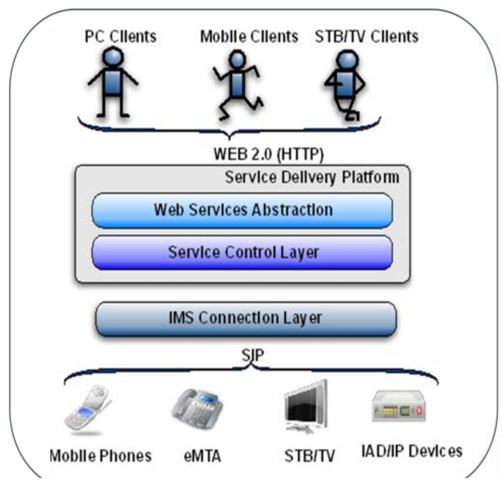

Figure 1 - Cable company's Service Delivery Platform

The Cable company required a Converged IMS Application Server that could handle SIP/ISC for session processing, diameter for rating, http for web services, and no downtime or failures. With services continually enhanced over the life of the cable/IMS network, the cable company required complete control over creating, integrating, and deploying the SIP Servlet building blocks. The cable company owned the services in order to control the content, enhancements and upgrades for their customers.

Lower Cost of Ownership

Fault Reliance and High Availability (HA) are not easy and taxes network resources with replicated hardware, and additional software to manage sessions across redundant systems. A redundant system also requires high speed internet connections; high speed replicated, and sophisticated processing to account for the transactional nature of communications. The overall performance of an HA system has a direct correlation to cost – the greater processing achieved in a shorter amount of time requires a smaller hardware/software footprint. If more hardware and software is required to process the sessions, the greater the overall cost.

A typical SIP connection requires approximately 13 transactional messages that must be sent and received in the right sequence in order to set up the session and tear down the session. The architecture must account and compensate for failures so as to preserve the integrity of the active sessions. All the computers and databases in both architectures are connected with a high speed network connection where messages are passed between systems to keep them synchronized. Since the architecture is built on networking between databases and computers, the overall performance is directly proportional to the performance of network or performance of the database – often measure in Milliseconds (1/100o of a second).

Data Grid is a shared in-memory set of data objects that provided low response time, high throughput, high scalability, and continuous reliability. Implementing a HA or Hot-Standby system in a Data Grid removes the local networking and the database performance overhead. In-Memory performance is often measured in nanoseconds (1/1,000,000 of a second), not milliseconds, that is, processing data on the grid is about 1,000 times faster. Since event processing is also built into Data Grid, events generated in a communications network maybe directly mapped into data objects which may trigger stateful processing on a session.

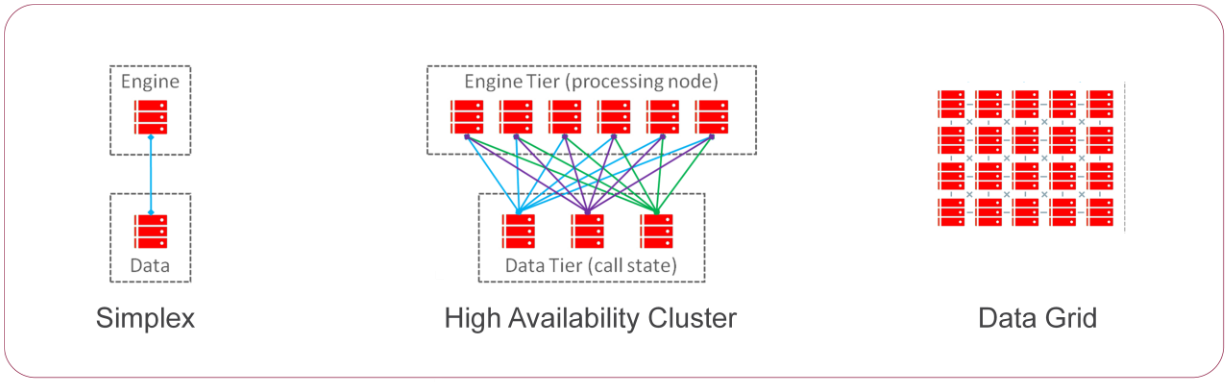

Figure 2 - Resilient configurations - Simplex, Highly Available, Data Grid

A Data Grid is designed for shared memory across application servers, which a CAS uses to back up data and process for failover and recovery or a Communications services may use multithreaded operation to access the data from different nodes. This not only speeds up recovery upon failure, but also simplifies highly complex services such as conference or collaboration across multiple communications elements. The Data Grid allows for the dynamic allocation or reduction of nodes in an active network, which fits nicely into a Virtualized Network or cloud architecture. The Data Grid scales the data capacity based on the network capacity. A communications service ported to a Data Grid performs on averages 3 times faster than a traditional HA system. This translates to roughly 3 times less hardware and software reducing the overall operational cost of network services.

For CSPs, consolidating footprint has a direct effect on operational expenditures. A very large wireless/fixed-line CSP was dealing with scaling issues, that is, growing in capacity and services without rebooting or taking the network or servers down, and without incurring higher operational costs. Instead of running back up systems in stand-by mode, they deployed the Oracle CAS in a High Availability (HA) cluster, were all systems were active. By running HA with highly efficient hardware, their hardware requires were reduced 5 to 1, by 80%, and their capacity of calls per second more than double. Overall, they dramatically reduced their hardware footprint reducing their maintenance costs, which also gave them more room to expand proportional to their business growth. By orchestrating the applications and scaling the clusters, they could accurately map their size of the clusters to the specific applications deployed. This gave them a direct association of operating costs to business growth.

Agile, Scalable, Orchestrated, Virtualized Network Function

The overall cost of a service is often derived from hardware where the service logic is executed. As hardware costs go up, the cost of the service goes up. If the service logic is not portable across hardware platforms, the network equipment market is limited based on the price of the hardware. Virtualization defines a clear separation between hardware and software, freeing the market for fair competition on both hardware and software.

"...we see a growing interest in virtualizing the Service Delivery Platform within a private cloud, particularly among large multinational operators looking to better leverage shared resources across a virtualized architecture."

SDP SOFTWARE AND SERVICES
ANNUAL WORLDWIDE AND REGIONAL MARKET SHARE, SIZE, AND FORECASTS
INFONETICS, FEBRUARY 9, 2015

Virtualizing the CAS offers:

- Reduced equipment costs and consolidating service infrastructure
- Decrease the time to market for new services by minimizing the hardware dependencies
- Isolate the management of CAS away from Hardware for services, subscribers, and applications.
- Rapid scaling up or down of services based on network demand.

A virtualized CAS may be deployed as a cloud element for enterprise computing or as a Virtualized Network Function (VNF) in a Network Function Virtualization (NFV).

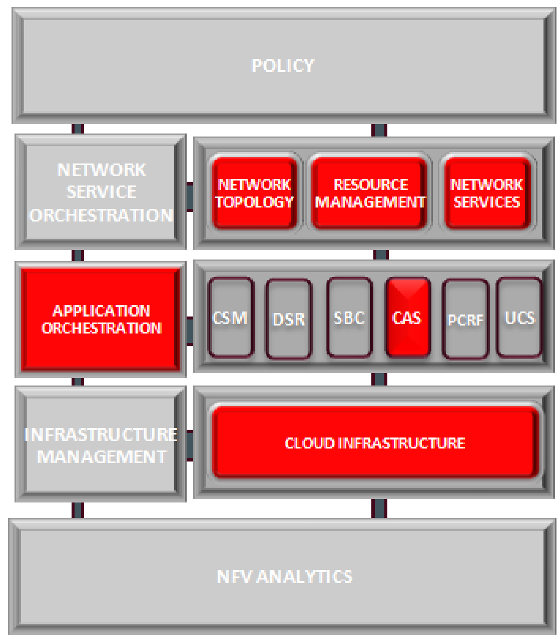

Figure 3 - CAS in an NFV Architecture

Deploying new services, taking on more customers, adding new features, all add capacity to the network, taxing the application servers. While most application servers can handle the growth of a few services and a few extra customers, CSPs face large scale growth of millions of customers and a several services added monthly.

One of the largest CSPs was faced with this kind of growth problem – a rapidly growing customer base with a demand for new services, at the same time attempting to migrate to an all-IP IMS network. Integrating into a Time Division Multiplexing (TDM) network with Voice over IP (VoIP) complicated by carrier interconnects and wireless interconnects, the CSP was facing a rapidly growing "cloud" based contact center business. Other integration issues included Automatic Call Distribution, Intelligent Voice Recognition, SIP Services, and Media Servers. The current network infrastructure could not keep up and they required quick changes to better position for rapid growth.

Facing major integration issues in the networks with services, the CSP deployed the highly scalable Oracle Communications Converged Application Server to integrate the applications and connect to the underlying IP network. Utilizing the orchestration features of the application router, defined by Java Specification 289, the CSP overcame issues of feature interaction and call routing, and readily introduced call center business logic as discrete applications within the container – see figure on the right.

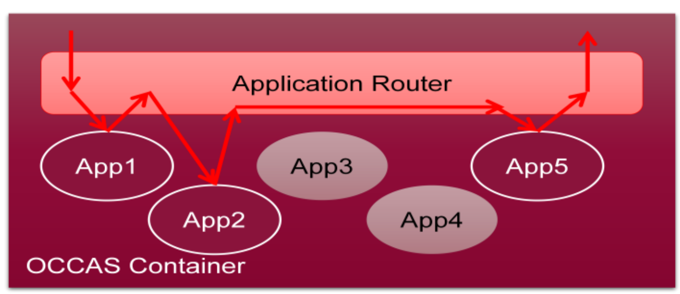

Figure 4 – Utilizing the orchestration features of the application router

Conclusion

The advances in a CAS are driven by the Web- Communications industry delivering an open, standards-based service creation and execution platform for IP-based communications applications. A CAS enables developing and delivering converged communications services rapidly and at a lower cost and thus enhances business agility and competitiveness. Aligning with the latest IT technologies, the CAS becomes far more agile with service creation and deployment utilizing the latest JEE standards. Performance and reliability improve dramatically with the underlying data grid. Virtualizing the CAS not only opens the market for a variety of hardware configurations, but also aligns the CAS for the future of NFV.

ORACLE

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

Advances in Communications Applications Server – Driving Innovation into Communications
January 2016