



QUALIFICATIONS PROFILE

- Results-oriented, innovative, hands-on senior executive with strong execution, strategic and long-range planning abilities and a track record of multiple successes delivering high-impact advanced analytic, large database and real time tick data products that drive revenue and grow market share, and improved product quality and development efficiency
- Diverse management background includes high performance server engineering, user experience, network operations, quality assurance, customer support and product management
- Proven leadership and team-building skills: managed a top-notch global development team of upto 125 in San Jose, Los Gatos, London, Paris, Tokyo, Beijing and Bangalore. Consistently displayed high pulse, engagement and management effectiveness scores. Managed groups smoothly through lean economic times and constant reorganizations without turnovers
- Excellent Communicator represented Relcy, Paypal, Thomson Reuters and Oracle at numerous customer visits and technical conferences. Well recognized in both industrial and Academic track in the area of very large databases. Also, hold about 30 patents in related technologies.
- Advisor to Vinod Khosla of Khosla ventures in the area of Big Data and databases.

EXPERIENCE

Relcy Inc (Sep' 2014 - Present) Chief Technology Officer

Brought on by the board to resuscitate Relcy Inc. - a mobile search company, after the founding engineers were hired by Google.

- Responsible for all of engineering and product direction.
 - Replatformed Relcy's search stack from almost concept to a highly scalable information retrieval platform.
 - Developed patented ideas for distributed crawling, grouping, enrichment, reconciliation and improved search query precision and recall by double (as measured through NDCG).
 - Grew the team from 3 to 15 engineers and stabilized operations (incl. data center ops), development and Q/A.

Paypal (an eBay company) (Mar' 2013- Aug' 2014) Global Head of Data Architecture & Infrastructure

Brought on board to head Data Architecture and Infra and bring in innovation and new technologies to the paypal data stack.

- Managed a team of 21 Architects & 65 + developers, responsible for data both OLTP & Warehousing, data solutions architecture and Big data infra and warehousing.
 - Provided strong technical leadership and serving as a role model for hands on data architecture. Attempted and solved a number of hard distributed data base problems which were outstanding for 5+ years which doubled #(payments) that could be processed through the paypal stack. This solution resulted in annual savings of 8+ million dollars in hardware.

- Responsible for providing innovative solutions to decrease cpu consumption for Teradata while processing 20% more data. Responsible for moving over expensive extract and load operations from Teradata to Hadoop, HBase, Hive and dual loading of data to improve availability and scalability of the data platform.
- Responsible for initiating the paypal graph database project to learn consumer behavior and create a recommendation engine for customers coming to paypal site.
- Delivered paypal credit data warehousing project at half the cost and ahead of time from concept.
- Key member of team tasked with securing IT systems posts eBay security breach.
- Member of committee responsible for evaluation of companies for M&A and Acquire.
- Member of committee drafted to recommend promotions across paypal to Director and Senior director.
- Instrumental in very successful contract negotiations with Teradata, Hortonworks and IBM for both hardware and software.

Thomson Reuters VP Content Architecture

Brought on board to revitalize innovation in the Enterprise division of Thomson Reuters Markets, a global software and content company dealing with financial market data and advanced analytics.

- Manage product development, Q/A, support and product management of the Reuters tick capture database, Tick History and Analytics with a staff of about 110+ employees/contractors across 5 locations.
 - Provided strong technical leadership and responsible for adding structure and process to Development teams resulting in faster release cycles and improved software reliability.
 - Responsible for introduction of big data technologies (Flume, Hadoop, Hive and HBase) for the Tick History Data Warehouse for tick and reference data collected off real time market data feeds.
 - Re-architecture of tick-database from a vertically scaled application running on expensive SMP boxes to distributed memory grid running on commodity hardware. Also, instrumental in moving historical tick data storage from a SAN to a clustered database using direct attached storage for performance and cost benefits
 - Evangelized benefits of basing tick databases off storage/compute grids internally and to key customers, resulting in almost doubling of revenues in a space of 14 months
 - Negotiated partnership agreements with Streambase, Portware, Matlab and Panopticon to expand the Data storage and Analytics platform
- Oversee development and QA for the reference data team of about 45 employees.
 - Review of reference data stack and capabilities of technology partners to go-to market faster
- Track technologies in the area of Big Data and middleware both in the industry and academic research, and fund projects of potential benefit to the company.
- Peer reviewer / part of Thomson Reuters CTO's team for all the tech stack in Markets division.

Oracle Corporation - (1997 - 2010)

Director Event Processing Products

Product development, management of distributed teams, strategy, acquisitions, architecture, implementation, team building, standards work, and product management of event processing middleware, products. Focused on bringing technical coherence and maximizing business proposition for Oracle in the Integration and SOA segment in particular. Member of Corporate Architecture Board.

- Managed acquisition and integration of the BEA event processing team and products into Oracle's middleware event processing stack.
- Managed Design, Development and Delivery of "Event Stream processor", powered by Oracle Continuous Query and Pattern recognition language from concept.
 - Domain specific server based off a OSGI/Spring DM based container that is designed for low latency, high throughput, scale out and high availability a first in the industry
 - Generic data management infrastructure for processing in-flight data before data is potentially stored to deliver results in near real-time
 - Rich event Stream processing engine that allows users to Aggregate, Correlate, Enrich and Detect Patterns in high speed streaming data in near real time
 - Uniform Declarative Framework which provides Data definition, Query language, Data Manipulation and System operations
 - Eclipse based Tooling to design event processing networks, construction of continuous queries, deployment and life cycle management of a CEP application
 - Console to monitor and manage CEP applications at run time, that also includes a component to do design time at run time
- Managed and evolved Oracle BAM acquired through the Peoplesoft acquisition.
 - Redesigned BAM core to use Oracle CEP as the event engine and Weblogic server as container
 - Added high-availability support to the event engine
 - Added Http pub/sub support for dashboards

Sr Manager Database Streams & Messaging

Given the challenging task of creating a generic information sharing product which is at the core of every distributed database service (replication, change data capture, logical standby, data pump) offered by Oracle RDBMS.

- Designed, Developed and delivered Oracle Data streams from concept

Oracle Data Streams enables the propagation and management of data, transactions and events in a data stream either within a database, or from one database to another. The stream routes published information to subscribed destinations. Streams contains three basic elements:

 - Capture

Oracle Streams supports capture of events (database changes due to DDLs & DMLs, and application generated messages) in the form of a logical change record (LCR) into a staging area.
 - Staging

Captured, events are placed in a staging area. The staging area is a queue that provides a service to store and manage captured events. Staging provides a holding area with security, as well as auditing and tracking of LCR data.
 - Propagation

If the subscriber is another staging area, the event is propagated to the other staging area, either within the same database or in a remote database, as appropriate.
 - Consumption

Events in a staging area are consumed by the apply engine, where the changes they represent are applied to a database, or are consumed by an application. Oracle Streams includes a flexible apply engine that allows use of a standard or custom apply function.

Sr. Manager - Queuing & Messaging, SQL Execution

- Designed developed and delivered a database integrated messaging and queuing product that allows the most sophisticated publish, subscribe and propagate applications to scale, and perform at requirements specified by Fortune 500 company architectures.
 - Key features are enterprise level scalable Point-to-Point, Point-to-Multipoint, Publish-Subscribe with support for Demand Driven Publication and Mutual Filtering
 - Built-in support for Rule based transformations
 - Structured Payloads with rich content, Retention and Message History for Auditing, Tracking, Security and Non-repudiation
 - Support for Rules including support for rules on aggregate, plsql functions and Directed Rules support through rule sets
 - Support for Internet Data Access and Propagation with XML Payloads
- Designed, Developed Delivered and supported a number of enhancements in the SQL Execution Layer in Oracle DB. Some key enhancements include:
 - Shared cursor layer for both Local & Distributed SQL
 - Serial reuse
 - Implementation of bundled execution in Oracle RDBMS, which is used by almost every layer in the DB
 - Supporting DDLs

Principal Member of Technical Staff

- Designed and developed parts of the best, high performance, ISV friendly C APIs for direct and distributed database access in the industry. The implementation with optimization was faster by a factor of 10 than the API that it replaced. Some of the salient features of these implementations are:
 - Using Bundling to reduce round trips
 - Piecewise access for large no. of rows
 - Dynamic Binds and Defines for handling bulk data

Consilium Inc. Mountain view, CA (1992-97)

Staff Engineer

- Primary architect tasked to migrate from a proprietary message bus solution to Corba 2.0. Designed, developed extensions using CORBA to include State transfer, Fault tolerance Load Balancing and guaranteed execution facilities.
- Designed and Implemented a Common Message bus API interoperable over UNIX and VMS. The API supports a rich features set such as:
 - Guaranteed executions
 - Spooling, Publish and Subscribe mechanism, Synchronous/Asynchronous messaging to groups of processes
 - Point to point messaging
 - Fault tolerant and redundant server computing, Process monitors, Process group watches, Load balancing and hot backup
 -

Sr.Engineer

- Designed and implemented Workstream Automated Material Handling system for Intel's FAB-V. Product allows station controllers to be served by Workstream Database, and integrates Intel's proprietary message bus S11 through gateways using VMS Mailboxes and UNIX sockets.

Digital Electronic Corporation, MA (1991-92)

Member of Technical Staff

- Technical lead for Motif Validation test suite (MVTs). MVTs is a suite of programs internally used by Ultrix Engineering group at DEC to test motif libraries before release and sign-off.
- Developed components of Digital's License management facility (LMF) for Ultrix. LMF is a network license scheme complete with licenses, distribution and usage reporting.

Shriram Industrial Enterprises Ltd, India (1986-91)

Manager

- Designed and implemented network abstraction layer above UDP. The abstraction provides reliability and recovery on an inherently unreliable medium.
- Implemented a memory allocator for DCM's proprietary flavor of UNIX. The system includes a powerful memory debugger and supports utilities that are used to investigate core leaks and memory violations.
- Technical lead in the systems division was responsible for operating system interface software that provided other developer's productivity tools like hash and balanced binary tree packages, string caching and error handling.

System Analyst

- Designed and developed a share accounting and public deposit's application for the DCM LTD corporate office.
- Developed core modules of a financial accountancy application.

Tata Institute of Fundamental Research (1983-85)

Research Assistant

- Worked as a research scholar in the Area of Algebraic Geometry. Contributed to some partial results relating to Jacobian Conjecture. Gave courses in:
 - Theory of Algebraic Surfaces
 - Category theory and spectral sequences
 - Algebraic Topology and theory of schemes
 - Algebraic number theory and Grassmanian Manifolds

EDUCATION

Indian Institute of Technology 1977-82

Kanpur, India

MS Major: Mathematics Concentration: Computer Science

Talks & Publications

- Key note address on Data clouds - Design & implementation (Debs 2013 Dallas)
- Server side Matlab integration with Tick Databases(Mathworks Chicago)
- Tick Databases (Market Microstructure conference Paris)
- Towards a Streaming Standard (VLDB Auckland 2008)
- Event Stream Processing (DEBS 2008 - Rome)
- Stream goes main stream (Panel Discussion Member - Sigmod 2007 - Beijing)
- Event Processing using Database technology (Sigmod - 2007 - Beijing)
- Grid Data Distribution at GGF9
- Best practices in Advanced Queuing
- Information sharing with the Oracle database at DEBS 2003

Patents Awarded

- 8676841 - Detection of recurring non-occurrences of events using pattern matching
- 8589436 - Techniques for performing regular expression-based pattern matching in data streams

- 8527458 - Logging framework for a data stream processing server
- 8498956 - Techniques for matching a certain class of regular expression-based patterns in data streams
- 8447744 - Extensibility platform using data cartridges
- 8397244 - Publisher flow control and bounded guaranteed delivery for message queues
- 8387076 - Standardized database connectivity support for an event processing server
- 8386466 - Log visualization tool for a data stream processing server
- 8352517 - Infrastructure for spilling pages to a persistent store
- 8326822 - Methods and systems for querying event streams using multiple event processors
- 8321450 - Standardized database connectivity support for an event processing server in an embedded context
- 8274115 - Hybrid orientation substrate with stress layer
- 8236678 - Tunable spacers for improved gapfill
- 8178417 - Method of forming shallow trench isolation structures for integrated circuits
- 8143651 - Nested and isolated transistors with reduced impedance difference
- 8145859 - Method and system for spilling from a queue to a persistent store
- 8103655 - Specifying a family of logics defining windows in data stream management systems
- 8053327 - Method of manufacture of an integrated circuit system with self-aligned isolation structures
- 8019747 - Facilitating flexible windows in data stream management systems
- 7984040 - Methods and systems for querying event streams using multiple event processors
- 7945540 - Method to create a partition-by time/tuple-based window in an event processing service
- 7932178 - Integrated circuit having a plurality of MOSFET devices
- 7818386 - Repeatable message streams for message queues in distributed systems
- 7779418 - Publisher flow control and bounded guaranteed delivery for message queues
- 7767577 - Nested and isolated transistors with reduced impedance difference
- 7761413 - Method of ensuring availability of event notification registrations of a database management system
- 7315863 - Method and apparatus for automatic notification of database events
- 7203706 - Buffered message queue architecture for database management systems with memory optimizations and "zero copy" buffered message queue
- 7185033 - Buffered message queue architecture for database management systems with unlimited buffered message queue with limited shared memory
- 7185034 - Buffered message queue architecture for database management systems with guaranteed at least once delivery
- 7181482 - Buffered message queue architecture for database management systems
- 7089265 - Database management system for implementing independent database actions in response to events of interest

Technical Summary

- Linux, Solaris 2.6/2.7, SUN-OS, HPUX, OSF/1, AIX, VMS, Windows-nt operating systems.
- C++, C, YACC, LEX, Lisp, csh, sh, awk, perl, pascal, fortran, XML languages.
- SOAP, X-Windows, Motif, DCE, UDP & TCP/IP, NFS, sccs client and server environments.
- Hadoop, HBase, Cassandra, Oracle 8.x.x, OCI, PL/SQL, Redis
- Streams, Advanced Queue, Isis, decmessageQ, Teknekron Information Bus, Orbix.

References

Available on Request