

CASE STUDY • PEOPLE MATTERS INFINI-T AWARDS INDIA 2026

One Vedanta, One Source of Truth: The Platform Other Platforms Plug Into

Vedanta Group (Vedanta Limited)

Infini-T Awards Category: People Platforms & Ecosystem Architecture



Executive Summary

~18,000 Employees across Pan-India and MENA operations	16 Group companies, all live on one core	6 Enterprise systems synchronised to the Darwinbox Hub	V-AIKYAM One Vedanta, one HCM: our own name for it
--	--	--	--

Vedanta is one of the world’s largest and most complex diversified natural resources groups: **16 group companies across seven sectors, aluminium, zinc, iron ore, steel, power, oil and gas and corporate, generating around US\$17 billion in revenue**, headquartered in India with operations spanning Pan-India and MENA. Our workforce runs past 100,000 people from white collar to blue collar, onroll and offroll, with **roughly 18,000 employees managed on the single HR core**.

Each of our entities had grown on its own HR footprint anchored in a legacy HCM, with employee data entered and re-entered by hand across payroll, learning and analytics tools. The risk was never a missing feature. **There were 16 different versions of the truth about the same workforce**.

The decision that defines this story is what we refused to frame it as. It was never “replace the previous HCM with Darwinbox.” It was to make Darwinbox **the single source of truth that SAP HANA and FICO, Disprz, Power BI, MMT, HGS payroll and Active Directory all had to sync with**, pushing hiring, transfer and separation events outward through open APIs and custom tables instead of letting each system keep its own copy of the same person. Integration depth and openness is why we chose **Darwinbox**: it could plug into our existing SAP backbone rather than force a rip-and-replace, with cost efficiency the second deciding factor.

We branded it **V-AIKYAM**, unity or oneness: one Vedanta, one system. Today **all 16 group companies are live on the single Darwinbox core**, feeding six connected systems around it. Most enterprises pick up a system of record and hope the rest of the stack falls in line. We engineered the fall-in-line.

What Group Scale Actually Means

This transformation did not happen inside one company. It happened across **16 group companies in seven sectors**, each with decades of its own HR practices, spanning white-collar corporate offices and blue-collar plant and mine operations across Pan-India and MENA. The starting landscape mixed legacy on-premises systems with newer cloud tools: **11 different systems across our group companies held scattered pieces of the same workforce**, which made a holistic view of the organization impossible. Getting one core to hold across that spread, sector by sector and entity by entity, is the scale claim of this submission.

100K+ People across the group, white collar to blue collar	7 Sectors: aluminium, zinc, iron ore, steel, power, oil & gas, corporate	US\$17bn In group revenue across the portfolio	11 → 1 Scattered systems consolidated into a single core
--	--	--	--

Getting one core to hold across that spread, sector by sector and entity by entity, is the scale claim of this submission.

Sixteen Versions of the Same Person

Before V-AIKYAM, we ran on SAP HANA and SAP FICO as the core HR and finance system, with learning, analytics, travel and other functions on separate, disconnected point tools. Each of our 16 entities had grown its own HR footprint on top of that backbone. The consequence was not that anything was missing. It was that the same employee existed in several systems at once, maintained by hand in each, and no one system could settle an argument about which record was right.

The cost was never only an integration cost. With data sitting in silos, leadership had **no group-level view of the workforce**, so people decisions could not be made proactively or promptly. Performance management ran **manually, and only for the leadership layer**, while each business unit worked to its own processes with little alignment between them. The problem we set ourselves was therefore wider than connecting systems: give the group one version of the truth, then make that truth usable, for every employee on a phone and for leadership on a dashboard.

Give the group one version of the truth, then make that truth usable, for every employee on a phone and for leadership on a dashboard.

Where Group Complexity Was Costing Our Business

The challenge we faced	The consequence it caused
Harmonising policies and user assignments across 16 entities	Inconsistent policy mapping and role assignment at entity level created duplicate records and rework, and forced a staggered, entity-by-entity go-live rather than a single group-wide cutover
Migrating the core system of record from legacy on-premises SAP to a SaaS-based Darwinbox core	This meant re-architecting who owned employee data, not just changing the interface. Change management had to address a technology shift and a mental one: from SAP owns HR data to Darwinbox owns HR data and SAP consumes it
Keeping SAP HANA and FICO, Disprz, Power BI, MMT, HGS payroll and Active Directory continuously in sync through push-based integration	Any gap in the event-based triggers for hiring, transfer and separation risked data drift across six systems, so integration governance and monitoring had to be built in from day one rather than bolted on afterward
No group-level visibility for leadership	Data silos across a mixed cloud and on-premise landscape made a holistic view of the organisation impossible, so workforce decisions were reactive rather than proactive, and leadership questions began with reconciling numbers instead of answering them
Performance management that was manual and limited to leadership	The appraisal process ran manually and covered only the leadership layer, so a group of our scale had no common digital performance rhythm connecting individual goals to business KPIs across entities
Standardising People SOPs across a diversified group spanning mining, metals and oil and gas, each with very different entity-level HR practices	The hardest part of the rollout was never the technology. It was getting 16 entities to agree on the common processes a true single-source-of-truth model demands
Driving group-wide leadership buy-in for a centralised system	Without visible sponsorship from group and entity CHROs, individual business units could have kept treating Darwinbox as one more tool instead of the system of record. Steerco-level engagement was required to convert buy-in into sustained adoption

One Record, Pushed Outward

The decision that shaped V-AIKYAM was architectural rather than competitive. We did not set out to displace anything that was already working. We set out to settle one question that sixteen entities had sixteen answers to: **who owns the truth about an employee**. Everything else followed from answering it.

We positioned Darwinbox as **strategic infrastructure, the hub of a platform-of-platforms**, rather than a stand-alone HR tool. The core holds the employee record. Every other system in our ecosystem, whatever its function, draws from what the core knows rather than maintaining a copy of its own.

Mechanically, that means employee lifecycle events (**hiring, transfer and separation**) are pushed outward from the Darwinbox core through open APIs and custom tables, event by event, as they happen. Nothing downstream must poll, request or reconcile. Getting there also meant consolidation: **11 different systems that had held scattered workforce data across our group companies were unified into the one core**, with biometric attendance devices and the surrounding enterprise systems integrated to it. Our north star from day one was to harmonise data maintenance and process across every entity, and we built the integration architecture to serve that vision rather than stitching it together after the fact.

That design is also why the decision went the way it did: **integration depth and openness**. Open APIs and a custom-table architecture meant the core could take its place inside our existing enterprise landscape without a rip-and-replace programme, with cost efficiency the second deciding factor.

Who owns the truth about an employee. Everything else followed from answering it.

One Employee, One ID

The same architecture settles identity. **Every employee is issued a single Darwinbox-native ID at the point of entry, and that ID persists as the sole identifier across their entire lifecycle, including moves between our group companies.** Internal mobility across sectors, aluminium to zinc to oil and gas to power to iron ore, becomes a lateral data update rather than a re-onboarding event: no duplicate records, and a continuous employee history regardless of which entity a person moves into.

Internal mobility across sectors becomes a lateral data update rather than a re-onboarding event: no duplicate records, and a continuous employee history.

Six Systems, One Record

Connected system	What it draws from the Darwinbox core
Finance and ERP	Employee master data synchronised to our finance and legacy core processes from the single record, with expense records processed in Darwinbox and pushed into SAP
Disprz	Learning, with populations and profiles kept current from the core rather than maintained in parallel
Power BI dashboards	Leadership analytics, the employee one-pager view and job rotation tracking, running on one workforce dataset instead of sixteen
MMT	Travel, driven from the same employee record
HGS	Payroll partner processing, running downstream from the core rather than from a separate master
Active Directory	Identity and access, provisioned and de-provisioned automatically from lifecycle events

Darwinbox carries the core HCM and all main modules, including Talent Management, Time Management, Recruitment and Onboarding. But the differentiation here is not the depth of any single module. **It is that sixteen entities and six enterprise systems now resolve to one employee record, and the core is what they all resolve to.**

What One Record Made Possible for People

A single source of truth is only worth building if it changes what people can do. Around the core, we rebuilt the daily experience of work across the group, and the benefits landed well beyond the integration layer.

What we built on the core	What it changed
Process harmonisation across the group	30+ HR processes harmonised across 16 group companies and seven sectors. - Employee lifecycle workflows such as confirmation and separation standardised. - One common performance management process for the entire group where each business had previously run its own
A single attendance and frontline framework	- One attendance system covering blue-collar and white-collar workers, integrated with biometric devices - Flexible shift management with auto-assignment, comprehensive rostering and auto-calculated overtime - Contractor performance tracked through custom scorecard forms on an integrated view of contract management
A mobile-first employee experience	100% employee and manager self-service on a mobile-first experience. - Recognition, our Vibe engagement page and a query redressal mechanism on the same platform, connecting a workforce that runs from corporate offices to plants and mines
Digital talent and performance for the group	- One digital PMS process for the whole group, run bi-annually with MBO-aligned KPIs, calibration and mobile-enabled reviews. 24x7 digital talent search on a single talent repository; V-Campus mentor-mentee programmes, IDPs for high-potential employees and successors for critical roles identified and managed on the system, with MSF/360 feedback live across entities

How Sixteen Entities Learned to Move as One

We staged the rollout deliberately rather than mandating it, and we treated governance and change management as workstreams in their own right, not as project hygiene. Each entity got a **SPOC enablement model**, a single point of contact owning adoption locally. A **central programme team** owned the overall rollout. And **monthly steering committee reviews with entity-level CHROs** tracked and discussed adoption numbers, which is what converted leadership buy-in into sustained adoption rather than a launch-day spike.

Governance also runs in the other direction, from the record back to the boardroom. With Darwinbox as the single source of truth, workforce data is pushed outward to **Power BI dashboards built for leadership-level visibility**, including job rotation tracking and a one-page employee view. At its sharpest, that becomes the **Chairman's dashboard**: a one-stop view of our people covering each employee's Vedanta career history and talent status, single-page organisation demographics such as gender diversity, average age and average tenure, the position status of key personnel against the business plan, and six-month workforce trends in hiring and attrition. The loop between operational HR data and strategic decision-making is closed: the numbers leadership steers by are the numbers the core maintains.

Change management is what carried it. The implementation approach was built to scale on a microservices architecture, go-lives were staged entity by entity, and policy harmonization ran ahead of technical rollout, so each entity adopted a process it had already agreed to rather than a tool it had been handed. The result: **all 16 group companies are live on Darwinbox as the core HR system**, across Pan-India and MENA, each feeding the same six connected systems.

The numbers leadership steers by are the numbers the core maintains.

The Critical Success Factors, as We Would Name Them

Critical success factor	How it worked in practice
Steerco cadence with entity CHROs	Monthly reviews tracked adoption numbers entity by entity, keeping sponsorship visible and sustained well past go-live
SPOC enablement model	One named owner of adoption inside each entity, close to local practice and accountable for it
Central programme team	A single team owning sequence, standards and integration governance across all 16 go-lives
Process before platform	Policy harmonisation ran as its own workstream ahead of technical go-live, so the core enforced processes the entities had already agreed
A governed data loop	The single source of truth feeds Power BI and the Chairman’s dashboard, closing the loop between HR operations and leadership decisions

What the Numbers Now Show

The architecture is real, the coverage is complete, and the flows through it are now measurable. The volumes below are what one record, pushed outward, looks like in production.

~4,000

People records moved between Darwinbox and SAP HC through onboarding and resignation events in the last year

~18,000

Expense records processed in Darwinbox and pushed to SAP through the Finance and ERP integration

30+

HR processes harmonized across 16 group companies

100%

Of the group on one digital PMS process, run bi-annually

100%

Employee and manager self-service, on a mobile-first experience

ONE

BI dashboard giving leadership full visibility on a single source of truth

Before and After, Measure by Measure

Measure	Before Darwinbox	After Darwinbox
Systems connected and data flowing	Disconnected, entity-siloed tools	6 systems integrated around the Darwinbox core
Entity coverage	Each entity on its own HR footprint anchored in SAP	All 16 entities live on one core, Pan-India and MENA
Adoption across integrated modules	Fragmented usage across separate tools	Live adoption confirmed across Disprz, MMT and SAP-linked workflows
Employee data ownership	Duplicated across SAP and point tools, maintained by hand	Core employee data maintained in Darwinbox and pushed outward on lifecycle events
Performance management	Manual, limited to the leadership layer	One digital PMS process for 100% of the group, run bi-annually, on mobile
Attendance and the frontline workforce	Fragmented tools split across blue- and white-collar workers	A single attendance system for the whole workforce, integrated with biometric devices
Leadership visibility	Data silos; no holistic view of the organization	Power BI and the Chairman’s dashboard running on one workforce dataset

Standardize the Process, Then Plug In the Tool

Phase one hit a problem that no integration could solve. Policy creation and user-assignment setup at entity level **ran into duplication issues**: the same employee or role occasionally ended up mapped incorrectly across entities. That forced a staggered, entity-by-entity go-live instead of the single group-wide cutover we had wanted.

The deeper lesson was harder. The hardest part was never integration engineering. It was **aligning People practices and SOPs across a diversified group** so a single central system could genuinely act as the source of truth. Mining, metals and oil and gas had each built their own HR practices over decades, and a single-source-of-truth model does not tolerate sixteen answers to the same question.

So, the sequence changed. We moved from “plug in the tool” to **“standardise the process, then plug in the tool”**, with policy harmonisation becoming its own workstream ahead of technical go live. That reordering is why the core can be trusted now, and it is the part of this story most worth copying.

Standardise the process, then plug in the tool.

Why We Believe This Stands Apart



A platform-of-platforms, not a core-HR rollout

We did not just adopt Darwinbox as a system of record. We architected it as the hub of a wider people and work-tech ecosystem, with SAP for finance and legacy core, Disprz for learning, Power BI for analytics, MMT for travel, HGS for payroll and Active Directory for identity built around it as spokes. The differentiation is the breadth and depth of ecosystem integration, not the depth of any single module.



We engineered the fall-in-line

Most enterprises pick a system of record and hope the rest of the stack follows. We made following a design requirement: event-based pushes on hiring, transfer and separation through open APIs and custom tables, so no downstream system needs its own copy of the same person. SAP consumes rather than competes.



Group scale, not single-entity scale

Sixteen entities, each with its own legacy processes, across mining, metals and oil and gas, spanning Pan-India and MENA. The hard part was the policy-onboarding effort needed to get all of them onto one central data model. Models of this kind are commonly proven at single-entity scale; very few have been proven at group scale.



Process first, then platform

The rollout reordered itself when it had to: policy harmonization became its own workstream ahead of technical go-live, because a single source of truth cannot be installed on top of sixteen different answers. That sequencing is the difference between a system of record and a system everyone actually records in.



One employee, one ID, for life

Every employee carries a single Darwinbox-native ID from the point of entry, and it stays theirs across the entire lifecycle. Internal mobility across sectors, aluminium, zinc, oil and gas, power, iron ore, is a lateral data update, not a re-onboarding event. In a sixteen-company group where careers routinely cross entity lines, an identity model with no duplicates and no history breaks is the difference between one workforce and sixteen adjacent ones.



From record to boardroom, a closed loop

The same record that runs daily HR operations feeds Power BI and the Chairman's dashboard: career history, talent status, demographics, hiring and attrition trends. Most groups run analytics as a periodic export exercise; ours runs as a live loop, so the numbers leadership steers by are, by construction, the numbers the core maintains.

From Data Pipes to an Intelligence Layer

The architecture will not stop at moving data. The next phase takes it **beyond data pipes into an intelligence layer**: parent-level policy intelligence held once at the group, entity-specific process extensions layered beneath it, AI-guided compliance, and configurable workflow agents acting on the record rather than merely reading it.

The innovation surface for that phase is already in place. **Darwinbox Studio, the platform's low-code / no-code extension layer**, gives us a plug-and-play API library to connect new tools without customising the core, and the ability to stand up new KPIs on demand as business priorities shift. Because every extension reads from and writes to the same single source of truth, the group can keep adding capability without ever re-opening the question V-AIKYAM settled: who owns the truth about an employee.

The group can keep adding capability without ever re-opening the question V-AIKYAM settled: who owns the truth about an employee.

From Our Leadership

The View from the V-AIKYAM Programme

“

V-AIKYAM was never a technology implementation. It was our commitment to creating one trusted people data foundation across Vedanta. By harmonizing people processes and establishing a single source of truth, we have laid the foundation for a digital, data-driven and employee-centric future.

”

Piyush Verdia • General Manager, Hindustan Zinc Limited

CONCLUSION

The Benchmark We Set

A diversified group does not have an HR data problem. It has sixteen of them. Mining, metals and oil and gas, each entity grown over decades on its own footprint, each anchored in the same HCM backbone and each maintaining its own copy of the same person by hand across payroll, learning, analytics and travel. Nothing was missing. Everything was duplicated. The question “how many people work here, and who are they” had **sixteen defensible answers**, which is the same as having none.

The move that made V-AIKYAM work was refusing the fight everyone expects. We never framed this as replacing SAP. We framed it as deciding who owns the truth about an employee, and then making every other system, SAP included, subscribe to it. **Hiring, transfer and separation events push outward from the core through open APIs and custom tables**, into SAP HANA and FICO, Disprz, Power BI, MMT, HGS payroll and Active Directory. Six systems, one record, no reconciliation. And one ID per employee, carried for life across every company in the group.

And the part that no integration diagram shows: it did not work until the processes agreed. Policy harmonisation had to become its own workstream, ahead of technical go-live, because sixteen entities cannot share a single source of truth while running sixteen sets of SOPs. **Standardise the process, then plug in the tool**. Every entity CHRO had to sign up to that, in monthly steercos, one entity at a time, until all 16 were live.

This is the standard we set: most enterprises pick a system of record and hope the rest of the stack falls in line. We engineered the fall-in-line and built the platform that other platforms plug into.