HP'S VISION FOR TOMORROW'S IT

"COMPUTE" ENABLES INFRASTRUCTURE TO BECOME MORE CONVERGED, WORKLOAD OPTIMIZED, SOFTWARE DEFINED, & CLOUD READY

EXECUTIVE SUMMARY

The forces of mobility, cloud, security, and big data are driving a sea change in the role of IT. Their combined impact is creating new business opportunities while simultaneously transforming the customer engagement model. Faced with this environment, today's traditional hardware-centric view of IT and slow pace of innovation cannot meet future business requirements in a world where latency and scale rule the day. IT must shift to a services-centric mindset and become a strategic business differentiator. To enable this, businesses will need to change how they approach their IT infrastructure in order to maximize business outcomes.

Hewlett-Packard has an approach for this new services-centric IT mindset that they call "Compute". HP's vision will help guide customers on the journey to redefine their approach to IT and better align their infrastructure to deliver the right business outcomes. **Compute is a vast pool of processing resources with the ability to think, store, and connect.** These resources can be located anywhere, scaled to any workload, and are always available. The right Compute for the right workload can accelerate IT service delivery, lower service cost, and fuel business growth. HP believes it can deliver the right level of Compute to support all phases of their customers' journey: from today's consolidation and convergence, through services-oriented IT, and eventually into a completely flexible environment that anticipates and provisions on-the-fly. To deliver IT as a service effectively and achieve differentiated business outcomes, more than just platforms need to change. Processes and people will need to evolve over time also, as IT shifts toward this more agile vision where **traditionally siloed infrastructure becomes more converged, workload optimized, software defined, and cloud ready**.

TODAY'S IT CANNOT SOLVE TOMORROW'S CHALLENGES

Traditional enterprise IT has evolved over time in response to business needs: from mainframe to minicomputers and then to industry-standard scale-out servers. But today's IT infrastructure and approaches cannot handle the growing needs of mobility, cloud, big data, and security which demand more agility and flexibility.

FIGURE 1: PROCESSING POWER CAN'T KEEP PACE

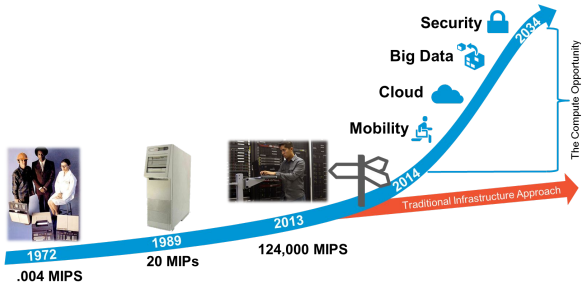

The “square peg” of business needs was always hard to fit into the “round hole” of IT, but companies did their best. Now, the velocity of change is making that peg even squarer while the hole becomes even rounder. IT is being driven by the need to reduce transactional latency and friction, especially as more commerce is being executed via mobile devices. What customers hold in their hands can impact their experience more than where they are physically located. In financial services, for instance, customers would queue at a bank teller or ATM for minutes, fully understanding the wait to execute the transaction. But online, a few seconds of latency while depositing or transferring money can severely impact customer experience. Such latency can drive a customer to switch vendors if their expectations are not being met. The competitive world of online and cloud services makes switching vendors easier than ever. If a business is not meeting customers’ needs, a different solution is just a few clicks away.

The velocity of business has changed dramatically. Years ago IT might have deployed only a handful of applications over the course of a year. Now IT is being pressured to deliver dozens of new applications per day. Today’s siloed IT is still loaded with rigid connectivity and manual processes that strain under the demand of all of these new applications. Management complexity abounds as IT staffs spend most of their time in reactive mode. The lack of granularity for changing and scaling platforms leads to a more brute force approach, which is not fine-tuned for the level of agility that new business models demand.

IT historically has been viewed as a cost instead of a differentiator. Thus, current infrastructures have been designed around reducing cost, physical footprint, and power consumption. This design philosophy leaves platforms ill equipped for the exploding amounts of data and processing that businesses now face. To compete effectively in the future, businesses need more efficient platforms. They need platforms that allow IT to focus on optimizing service delivery. IT then needs to focus on lowering the cost, reducing the delivery time, and increasing the business value of these services.

Comparing the world of IT to the US auto industry shows some interesting parallels. In the years between 1965 and 1975, vehicles had small incremental changes each year, similar to the way new functionality has come to server platforms for the last decade. But the oil embargo of 1973 completely changed the dynamics of the industry and caught US manufacturers off-guard. The market suddenly shifted. Japanese manufacturers, who had been watching and planning for disruption, sprang into action. These Japanese manufacturers anticipated changing trends in commuting times, fuel efficiency, and space. They redefined the market by anticipating change and delivering on these new customer needs. Meanwhile, some US automakers collapsed while others struggled through consolidation.

Similar disruption is happening to Enterprise IT. When markets are disrupted by trends like big data and mobility, the status quo stops working, and new enterprise competitors change the dynamics of a vertical industry. Because of today's market forces, the IT path needs to change. Companies that can respond quickly will succeed.

HP'S COMPUTE VISION AND APPROACH

HP's long-term vision for IT is called "Compute". Compute is a noun—not to be confused with the lowercase verb "compute". Because HP's Compute vision includes so much more than just physical processing, it surpasses the limited definition of the verb compute. In HP's Compute vision, IT infrastructures morph from physical systems into **flexible resource pools driven by virtual overlay states**. Allocated resources allow the system to scale up or down fluidly, based on business rules or environmental triggers. Beyond just processing, Compute transforms traditionally siloed, manual infrastructure, enabling infrastructure to become more converged, workload optimized, software defined, and cloud ready. Most importantly, Compute enables the infrastructure to deliver differentiated business outcomes both now and in the future.

Compute is a journey that can be divided into three phases: Optimized Compute, Predictive Compute, and Autonomic Compute.

1. **Optimized Compute** brings efficiency and productivity
2. **Predictive Compute** increases flexibility and resource utilization to enable truly services-oriented and cloud-like delivery
3. **Autonomic Compute** infrastructure scales to new levels and stays in lockstep with business needs

FIGURE 2: THREE PHASES OF HP COMPUTE

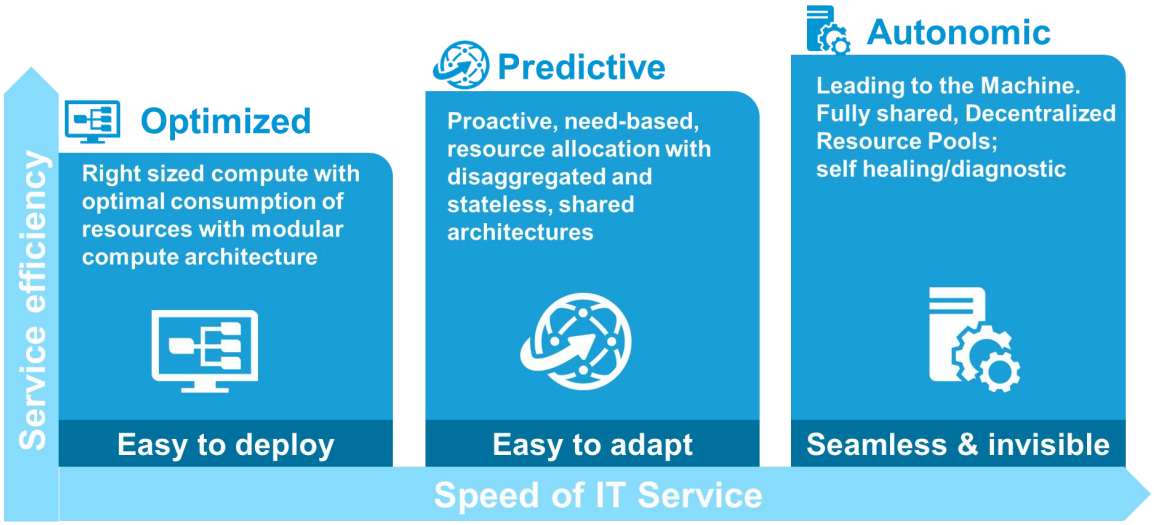

The technologies to enable the first phase of Compute are available today. Over time, additional technologies will become available to enable the future phases. The Compute vision encompasses all systems: from the cloud-based “systems-of-engagement” that are the direct interaction point with customers today, into the mission-critical back end “systems-of-record”, and all the way through to the “systems-of-analysis” that drive business insight.

Over the years, investment in systems-of-record traditionally focused on efficiency and managing costs. Customers were launching one or two new large applications per year, so the old development model worked. But business needs are changing. Now customers are launching dozens of smaller applications at a rapid pace in the systems-of-engagement and systems-of-analysis spaces. Because these smaller apps still require robust logistical provisioning and security, the old development model is breaking down. IT staff and traditional infrastructure cannot keep pace. Enterprises need a different IT model unbound by the traditional development approach and culture of the prior 20 years.

Compute investments require a holistic view across all these systems, as IT service delivery typically encompasses multiple systems to handle business needs seamlessly and effectively.

FIGURE 3: HOLISTIC SERVICE DELIVERY

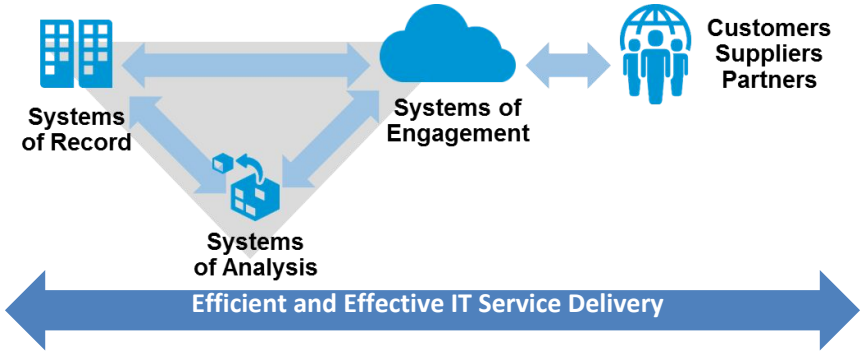

With today's traditional infrastructure, several different approaches are used to address these systems. In the future, the single Compute approach can deliver the right level of Compute that could bridge all of these systems. But it is far from a "one size fits all" approach, instead focusing on the right level Compute for each service.

PHASE ONE: OPTIMIZED COMPUTE

In this first phase, infrastructure begins to blend as the lines blur between platform & services, with an emphasis on reducing deployment & management costs while also improving businesses agility. The Optimized Compute phase focuses on breaking down silos, consolidating systems, and streamlining operations so that customers can reduce cost and increase efficiency. The key aspect of this phase is aligning the right platforms to the right processing tasks. Blades and consolidated infrastructures, for instance, play a valuable part in today's optimization efforts. These systems will continue to have a place in the future for back office applications but not necessarily for the systems-of-engagement and systems-of-analysis that are changing so rapidly today. These quickly shifting workloads are better suited for density-optimized and hyperscale platforms that can scale up rapidly.

In Optimized Compute, the lines between server, storage, and networking begin to blur. Converged IT deployments drive cross-functionality for both platforms and processes. East/west network traffic is growing and storage is becoming distributed; breaking down the silos and handling more functions inside the chassis or rack reduces communications latency. Along with breaking down the system silos, businesses will also need to break down the functional silos between their IT teams.

System granularity is still at the platform level, so resources are added and managed by the platform. Deployment time shrinks, allowing IT services to be spun up faster, so businesses can reduce the ROI breakeven point and drive greater profitability

earlier in the business cycle. As IT becomes services-focused, they can begin the next step towards Predictive Compute.

PHASE TWO: PREDICTIVE COMPUTE

The second phase of Compute is where a services-led IT enables proactive services deployment by anticipating the requirements of the business ahead of actual needs. In the Predictive Compute phase, two very distinct things occur.

- Deployment and management of IT resources shift from **workload**-focused to **workflow**-focused. Virtual overlays connect services as resource definition moves from hardware-defined to software-defined. This shift provides better insight into what **could** be occurring (*i.e.*, predicting), so systems can better align resources with needs.
- **Resource pools become more granular, down to the subsystem level**, for more flexibility and finer tuning, allowing systems to better scale up.

With an almost fully abstracted overlay, the underlying hardware now has the intelligence to interact with the orchestration and management layer, providing the data needed for predictability. Physical definition gives way to statelessness. Disaggregation of systems becomes the mechanism for delivering on-demand resource deployment. The integration of multiple disparate data sources (like environmental, social media, sensor data, and third party federated data) assists in sharpening the insight and predictability of the service's needs.

Stepping from the Predictive Compute phase to the final phase does have its prerequisites as well. One cannot move to the last phase until **all** of the underlying systems are fully virtualized and software-defined. IT alignment at this point is fully services-focused, and systems are automated to drive their own Compute levels.

PHASE THREE: AUTONOMIC COMPUTE

In the Autonomic Compute phase, **flexible resource pools are granular down to the component level**. The system is highly resilient and self-healing. The business is now fully virtualized and fully services-led. Services can now spawn other dependent services, spinning up or down resources based on need. **The system is now autonomic, with the ability to self-define, self-orchestrate, and self-heal.** Because so much of the infrastructure is software-defined, management is more like handling code than hardware. The IT organization has moved from systems managers to

services template designers who focus on the business instead of managing the ratios and resources between servers, storage, and networking.

With the highest level of component granularity, hot-pluggable modular components can be added to resource pools on the fly, thus removing the need to take systems offline to address capacity. The system sees trends, watches business processes, and dynamically allocates resources (irrespective of where the physical components reside)—all in real time to ensure that service levels are maintained or exceeded. System expansion is maximized, because it is at the most granular level and requires far less over-commitment of resources. This granularity has a tremendous savings on capital equipment.

As a completely services-based infrastructure, business agility is maximized, and businesses can move at top speed to capitalize not only on macro trends but even micro trends that arise quickly in the market.

THE HUMAN & PROCESS IMPACT

Much has been said about how systems will need to change in order to become more complementary for a services-oriented IT environment, but we cannot ignore the potential impact to both the IT workers and the underlying processes.

Throughout the Compute journey, IT staff will see their roles change. They will move from reactive hardware support to become more integrated into the business processes, helping align business needs with the new services orientation. **Today's hardware specialization disappears, as IT workers become proactive business partners instead of reactive silo protectors**.

The manual processes and hands-on configuration of today's world give way to more automation, as systems become more capable of running themselves. When provisioning and deployment can be spawned from a service instance and resources can be automatically allocated or removed based on business needs, then peak efficiency is achieved, and the cost of managing IT operations is dramatically reduced.

COMPUTE PHASE COMPARISON

Table 1 illustrates the differences between the phases of Compute.

TABLE 1: COMPUTE PHASE COMPARISON

Phase	1. Optimized Compute	2. Predictive Compute	3. Autonomic Compute
Goal	Converged deployment increments make it easier to react to changes	Align pools of software-defined resources to anticipate changes	Real-time automated response to changes with hands-off management
Description	Right sized compute with optimal consumption of converged resources	Proactive, need-based, disaggregated resource allocation using stateless "RESTful" architectures	Fully shared, decentralized resource pools, self-orchestrating, & self-healing
Resource Granularity	Discrete servers partitioned to match general workload profiles	Subsystems configured as virtual servers optimized for specific workloads	Components dynamically configured as virtual servers optimized for individual applications
Action	Begin the process of automation & orchestration, implement baseline private cloud technologies	Understand self-service provisioning & cloud bursting from private to offsite managed or public cloud (*i.e.* hybrid cloud)	Stay up-to-date on the latest software-defined infrastructure, cloud framework, & orchestration capabilities (*e.g.,* HP Helion)
Business Outcomes	Baseline abstraction between platform & services will reduce deployment & management costs & improve businesses agility	Services-led IT enables proactive services deployment by focusing on anticipating requirements of the business ahead of actual need	Autonomic services-led IT enables focus on optimally serving business needs instead of focusing on systems & software architecture

CALL TO ACTION

If we look back at our auto industry analogy through today's lens, we see that new business models are once again impacting the market. With fractional ownership (ZipCar, Car2Go, *etc*.) and ride-sharing services (Uber, Lyft, *etc*.), many young buyers are foregoing car purchases altogether. This disruption is creating tremendous upheaval as manufacturers scramble to keep pace with unanticipated innovations. The disruption goes beyond "hardware" manufacturers to transport services as well. Customers were once content to stand and wait for the next taxi to drive by. Now customers want control. They want to be able to see the car, know the price, know the driver, and conduct the transaction from the palm of their hand. Drivers were once content to prowl the streets for fares, but now drivers want control too. They want a verified customer, best-path routing, and seamless billing/payment.

The parallels with the IT industry are instructive, and they also demonstrate a use case. These new transport activities require a tremendous amount of IT capacity that can scale up or down instantly to meet demand. Execution speed is the new currency, and latency is the casualty of this new way of doing business. New business models are "breaking the mold" and forcing companies to take a very different look at their IT; they need to innovate or get left behind. New business models come from almost nowhere, imposing tremendous real-time IT tasks: tracking assets and customers, matching needs in the most efficient manner, balancing workloads across a disparate set of variables as well as instantly billing, and notifying all of the parties involved. Today's IT architectures just aren't designed to deliver all of these real-time functions. Something needs to change.

Many IT vendors see a gradual transition from the current siloed systems to a fully virtualized, services-focused IT future where disaggregated pools of resources can be provisioned dynamically to match the exact business needs. This future state vision promises to reduce the cost of computing while simultaneously boosting its ability to scale up and down based on needs. While other IT vendors may foresee a similar future, HP clearly has two advantages in making it happen. First, it owns a large share of the market, making HP an efficient starting point for many customers as they embark on this journey. Second, HP has one of the widest portfolios, allowing customers to evolve along their Compute journey without requiring a wholesale infrastructure change.

HP has articulated an evolving vision in a phased approach. As new technologies are available and business processes mature, HP's vision gives customers a framework for understanding and moving towards more automated, efficient, and cost-effective IT.

We recommend that customers take the first step by engaging HP in a discussion of their long-term strategy to see how HP's Compute vision can help move their IT to a more agile and cost-effective enabler for tomorrow's business.

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHOR
John Fruehe, Senior Analyst at Moor Insights & Strategy

REVIEW
Patrick Moorhead, Founder, President, & Principal Analyst at Moor Insights & Strategy
Paul Teich, CTO & Senior Analyst at Moor Insights & Strategy

EDITOR
Scott McCutcheon, Director of Research at Moor Insights & Strategy

INQUIRIES
Please contact us if you would like to discuss this report, and Moor Insights & Strategy will promptly respond.

CITATIONS
This note or paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

DISCLOSURES
Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER
The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.