Internet Distributed Computing

The Intersection of Web Services, Peer-to-Peer, and Grid Computing

David Barkai
Distributed Systems Labs
Corporate Technology Group
Intel Corporation

September 5, 2002

intel

Copyright © 2002 Intel Corporation.

Outline

- Why "Internet Distributed Computing" (IDC)
- What is IDC
- Technology Pillars
- Gaps and Challenges
- Future Applications

intel Copyright © 2002 Intel Corporation.

The Vision

- A computing world with
  - Sharing networked resources
    - Through dynamic pooling, aggregation, and load balancing
  - Automation enabled by
    - Reliable access to resources
    - Proactive and responsive processes and agents
    - Adaptive computing sensitive to changes of bandwidth and presence
  - Ultimate mobility with seamless and un-tethered access to computing resources in the environment

Copyright © 2002 Intel Corporation.

Motivation: Foundation for Evolving Computing Environments

- Proactive computing
  - react to real-world stimuli
  - initiate actions based on user's context (intent, activity, location)
- Pervasive/ubiquitous computing
  - grand scale: Internet of things (devices, sensors, objects)
    - Mobile, intermittent connectivity
  - compute resources all around in the environment
    - Islands of "smart spaces"
    - Users augment capabilities from the environment

Copyright © 2002 Intel Corporation.

Goals of Internet Distributed Computing (IDC)

- Accelerate distributed computing on the Internet, to increase business and consumer end-user value
- Provide a foundation for pervasive computing from a small-scale Personal Area Network (PAN) to virtual, planetary grids.

intel

Copyright © 2002 Intel Corporation.

Foundational Computing Domains

- Web Services
  - Standard protocols for service oriented architecture
  - Impetus in business process
- Grid Computing
  - Infrastructure services for sharing distributed resources
  - Origins in high performance computing community
- Peer-to-Peer Computing
  - Collaboration and computing at the edges of the net
  - Focused on client-side of network computing

intel Copyright © 2002 Intel Corporation.

Web Services: Protocols

- Standard, platform-independent, component-based distributed computing paradigm
  - for data exchange, capability descriptions, publishing and discovery mechanisms
  - Internet-centric, open, scalable, distributed components
- Fundamental standards/protocols for IDC
  - [OGSA*] SOAP for communication
  - [OGSA] WSDL and WS-Inspection for describing and locating services/resources
  - [desired] UDDI-like discovery mechanism

* "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration," I. Foster et al, June, 2002

SOAP: Simple Object Access Protocol
WSDL: Web Services Description Language
UDDI: Universal Description, Discovery, and Integration
OGSA: Open Grid Services Architecture

intel Copyright © 2002 Intel Corporation.

Grid Computing: Services

- Infrastructure and basic services for large-scale resource sharing
- Virtualization through focus on interoperable services
  - Services = cycles, storage, networks, databases, files,..
- Basic services
  - Resource allocation, resource management
  - Information discovery service
  - Security infrastructure
- Service semantics

intel Copyright © 2002 Intel Corporation.

P2P: Action to the Edge

- Peer-to-Peer technologies to enable PC/device to be a grid node and to offer services
  - In particular, naming mechanisms to deal with NAT/Firewall traversal (that applies to most users)
- Devices at the edge are the focus for mobility in IDC
- Foundation for mobile discovery and services
- Extend classes of applications with Collaboration and Content Sharing

Copyright © 2002 Intel Corporation.

Synergy with Trends and Drivers

- Web Services
  - Basis for common protocols, app-to-app communication
- Grid Computing
  - share and amplify computing resources, access to unique data
- Peer-to-Peer Computing
  - Direct communications for ad hoc comm and collaboration
- Autonomic computing
  - automated operation, self-management
- Utility computing
  - Efficient use of computing resources through resource pooling and load balancing, dynamic allocation/provisioning of servers, storage

intel Copyright © 2002 Intel Corporation.

A Unifying Computing Platform

The Internet as a distributed computing platform unified by common aspects of Web Services, Peer-to-Peer, and Grid computing, while also preserving the best of their unique, problem-solving aspects.

Copyright © 2002 Intel Corporation.

Converged Middleware: The Building Blocks

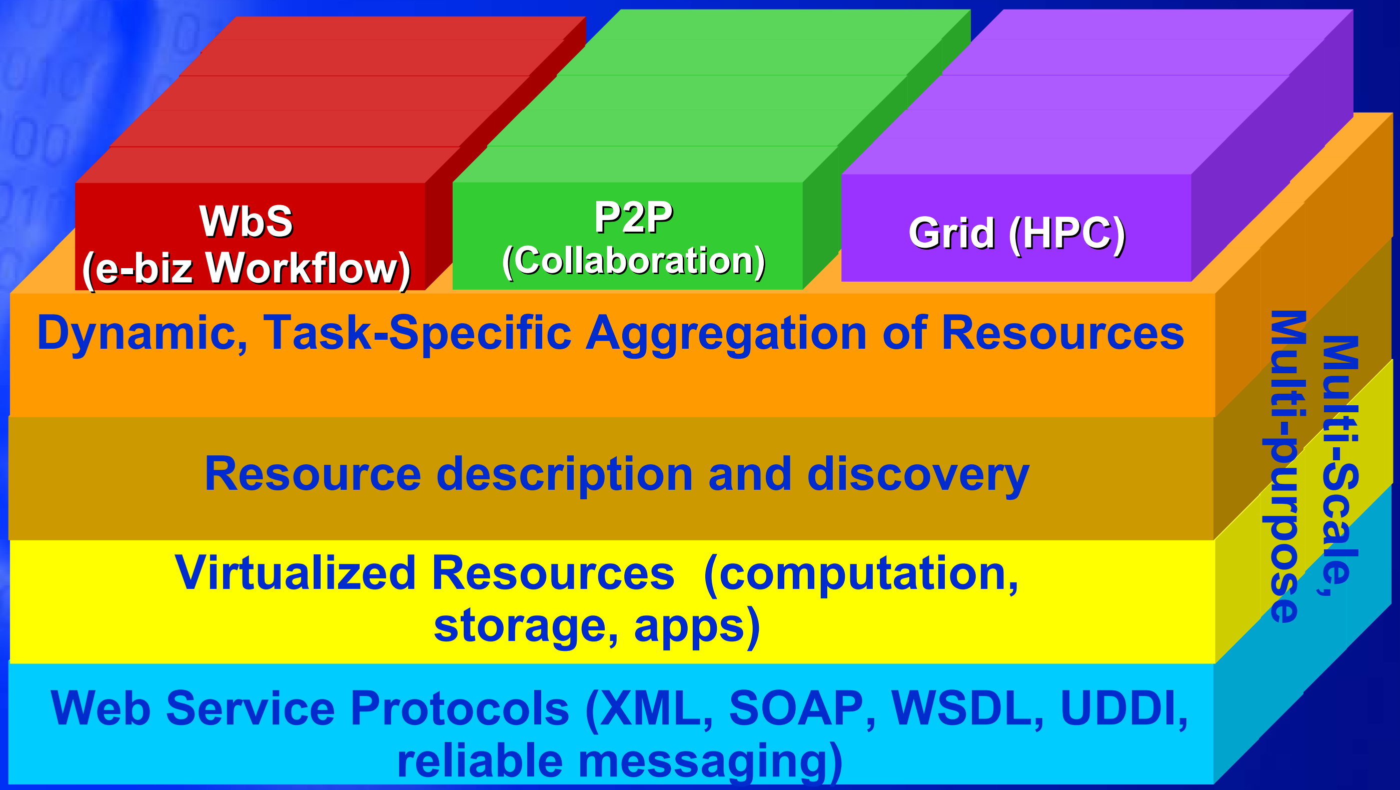

Apps

IDC =

G + P2P

+ WbS

Copyright © 2002 Intel Corporation.

IDC App Functional Stack

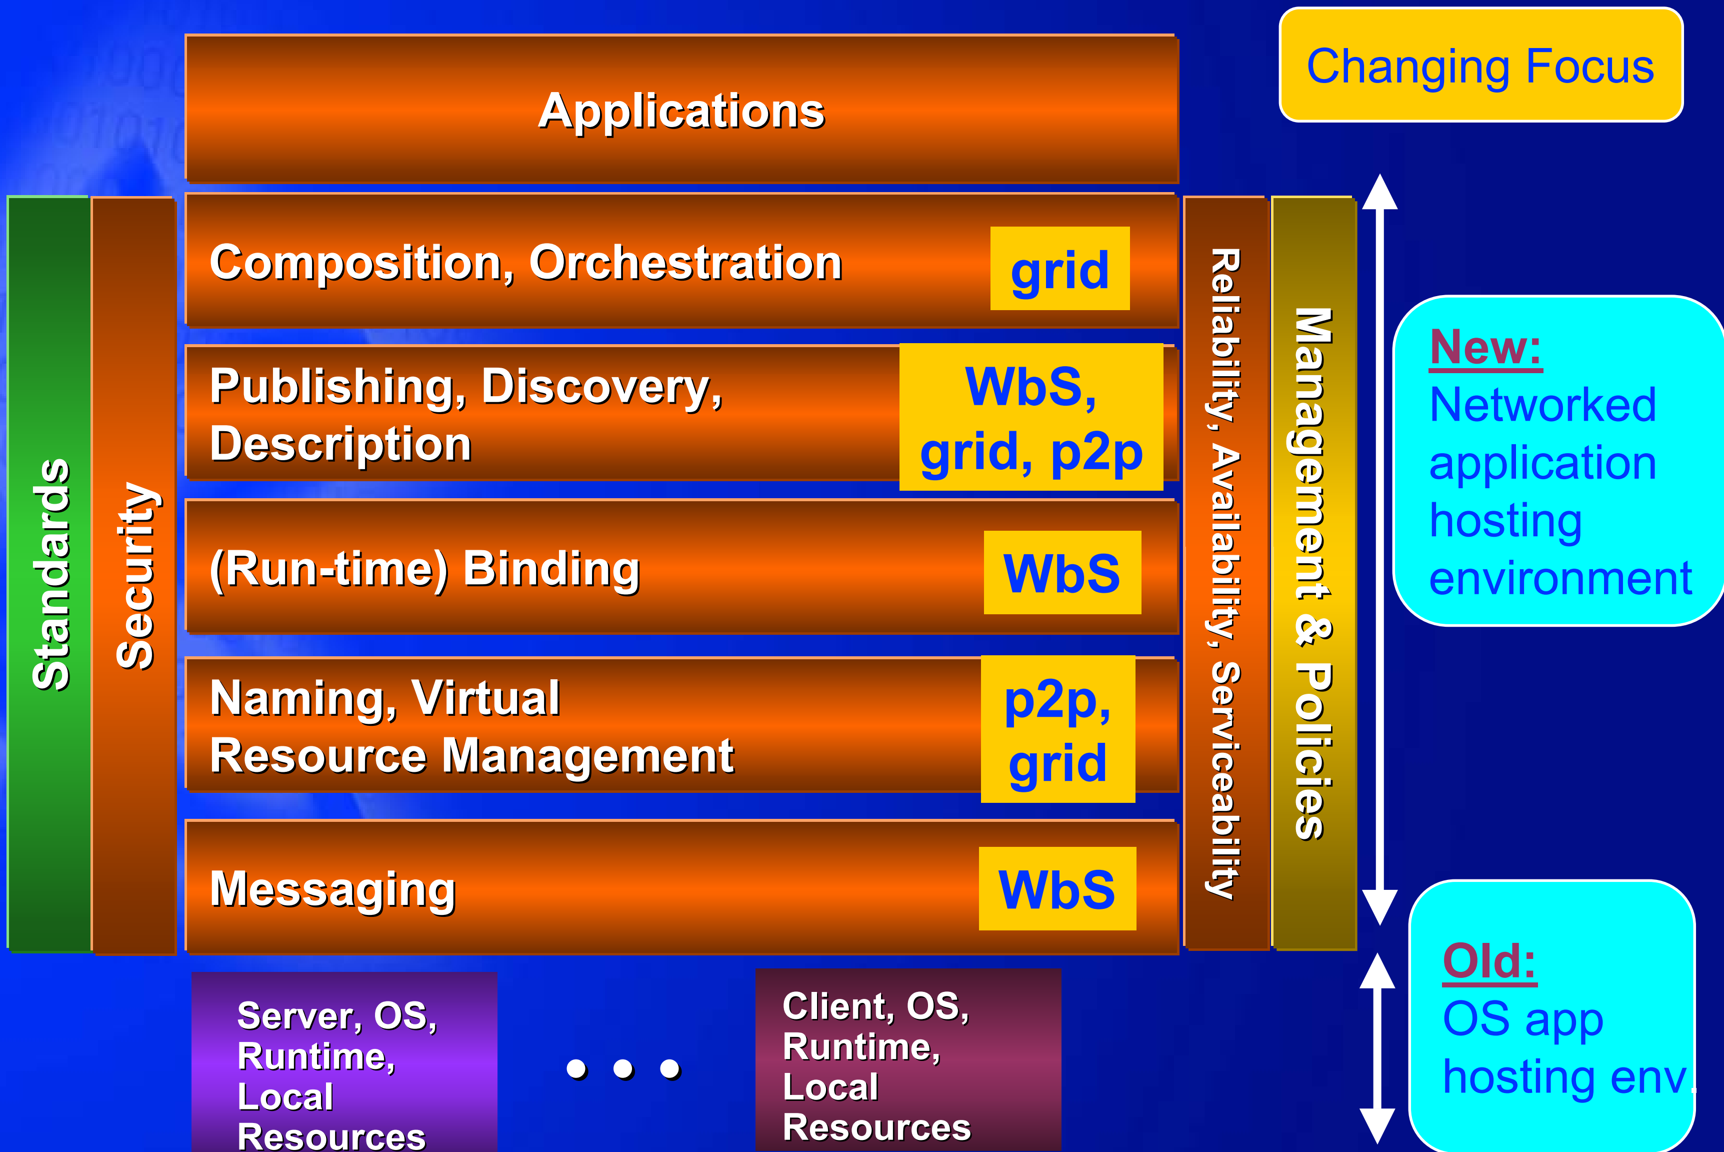

intel
Copyright © 2002 Intel Corporation.

The Benefits from IDC

- Dynamic aggregation of resources for computing on demand.
- Reusable SW and associated developer benefit.
- Foundation for proactive and pervasive computing.
- Reduce manageability cost and complexity.
- More efficient use of computing resources.
- Reliability, availability and scalability.

intel Copyright © 2002 Intel Corporation.

Realizing the Vision

- Define a common building-block architecture that scales from PAN to global grid
- Apply research and construct concept proofs of IDC fundamentals
- Work with key Standards/Industry groups,
  - W3C (SOAP), OASIS (UDDI), WS-I, GGF (P2PWG)

Copyright © 2002 Intel Corporation.

Elements of the IDC Model

- The Internet as a modular service-oriented computing platform
- Dynamic, task-driven configuration of computing elements
- A universe of globally connected services utilizing a vast pool of resources on all devices
- Dynamic discovery and combining of components into task-specific functional groupings
  - aggregate, augment device capability
  - capacity on demand, reusable: return to pool when task completed
- Multiple scales: proximity-area to global grids

Copyright © 2002 Intel Corporation.

Fundamental Principles

To realize the vision (scale from PAN to global grid):

- Virtualization
  - Of resources: cycles, storage, UI
- Automation
  - Of dynamic configuration and runtime binding
    - through resource discovery, ad-hoc networking, platform independence
- Modularity
  - For aggregation and orchestration of resources: pooling, dispatching, synchronization
- Trust
  - Through security, authentication - by ownership, proximity, organizational boundaries, intranet, Internet

Copyright © 2002 Intel Corporation.

Virtualization

- Virtualization through use of functional software components
- Apply to services and hardware elements
- Scalable to a planetary-size pool of composable resources
  - Connected via networks of networks
  - Described, accessed via standard web services protocols, directories, and conventions

Copyright © 2002 Intel Corporation.

Dynamic Binding

- Combination of discovery and binding mechanisms
- Application portability
  - Through decoupling of application design and execution from system configuration and physical connectivity
- Decouple development of uses from the provision of the service
- Necessary for ad-hoc and self-configuring networks
- Quality of Service
  - Improved reliability
  - Load balancing

Copyright © 2002 Intel Corporation.

Aggregation, Orchestration

- Dynamically aggregate the exact collection needed for task at hand – central to IDC
- Apps need to state requirements/hints
- Key element to vision of pervasive computing
- Orchestration: control and management of aggregated resources
  - Enables collective progress
  - Coordination and collating via synchronization, communication functionalities

intel Copyright © 2002 Intel Corporation.

Security, Authentication

- Goal is the protection and trust for both host and consumer/guest app
- Robust IDC allows crossing distinct security domains and varying levels of trust
  - Must not violate local policies (not a lowest common denominator approach)
- Single sign-on across aggregated collection of resources
  - Desirable, if not necessary

Copyright © 2002 Intel Corporation.

Design Implications

- Edge processing
  - Proximity to user – access to more devices, optimize bandwidth requirements
- Hybrid Connectivity – wired & wireless
  - WAN and broadband short-range wireless
  - Aggregate bandwidth at the edges
  - Ad-hoc networking
    - Discovery & self-configuring require device intelligence and autonomy
- Data+Communication-centric model
  - Majority of computations support access to content and connectivity
- Intermittent connectivity
  - Mobile devices move in and out of range, power down
  - Manage through caching, data and service proxies
- User-Environment integration
  - Users carry devices, but augment capabilities from the environment

intel

Copyright © 2002 Intel Corporation.

The Intersection, Today

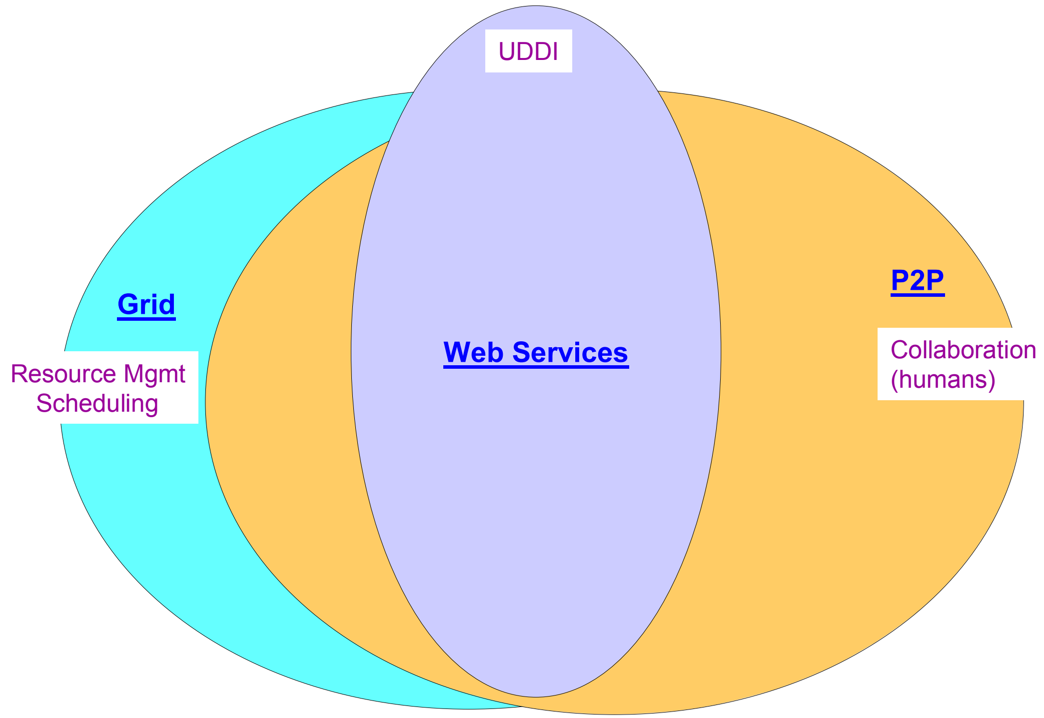

Some areas of divergence

intel
Copyright © 2002 Intel Corporation.

A Vision for the Future

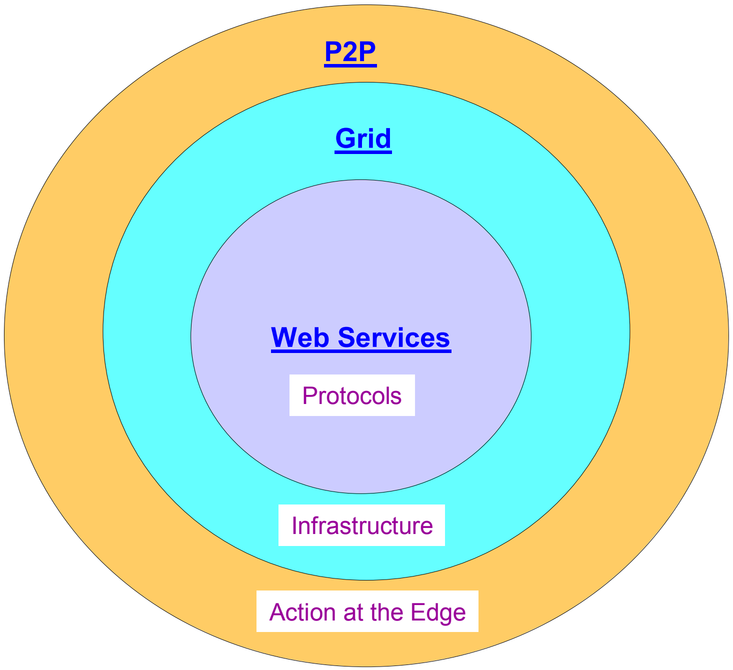

Foundation for a truly interoperable distributed computing

Copyright © 2002 Intel Corporation.

The Hard Problems

- Interoperability for mobility and access to environment
- Trust for sharing resources in the “cloud” and on own devices
- Connectivity for access, bandwidth, mobility, intermittency
- Virtualization of all resources
- Device Intelligence for automation, proactive, context-based computing

Manage complexity of dynamic, scaled-up, environment

intel

Copyright © 2002 Intel Corporation.

Missing from Grid+WbS

- PC/Client as a provider of services and as a full-fledge node on a grid
  - Getting to/from in presence of NAT and firewalls
  - Preserving local autonomy
  - Availability, allowing for users' whims and needs
- Allowance for Mobility
  - Location context awareness
  - Access to data when device is resource-constrained

Let's talk about Mobility this time …

intel
Copyright © 2002 Intel Corporation.

Mobility – The Missing Domain

- Current web services, grid computing, and most P2P are designed for static configuration
  - Albeit allowing for intermittent connectivity (some P2P)
- Challenge: Include mobile devices and users
  - Both as providers and consumers of services and resources

intel
Copyright © 2002 Intel Corporation.

Mobility – Embodiment of Challenges

- Mobility and rate of change
  - permanent, nomadic, moving
- Connectivity
  - full, intermittent, none
- Spectrum of trust/security
  - untrusted, mix, trusted
- Available internal resources
  - power, CPU, memory, etc
- Available external resources/infrastructure

Mobility extends IDC requirements for supporting dynamic and ad hoc usage models.

Copyright © 2002 Intel Corporation.

Handling Intermittent Connectivity

- Support for asynchronous communications, reliable messaging and caching.
  - Masks mobility (e.g. device on/off, here/not here).
  - Handling intermittent and unreliable connectivity.

Copyright © 2002 Intel Corporation.

Dynamic, Adaptive Security

- Provide automatic, dynamic, telescoping security.
  - Security with reduced user burden.

Copyright © 2002 Intel Corporation.

Sensitivity to Limited Resource

- Sensitivity to resource-constrained devices and communication channels, modulo Moore's Law.
  - E.g. power, network traffic, memory footprint, etc.

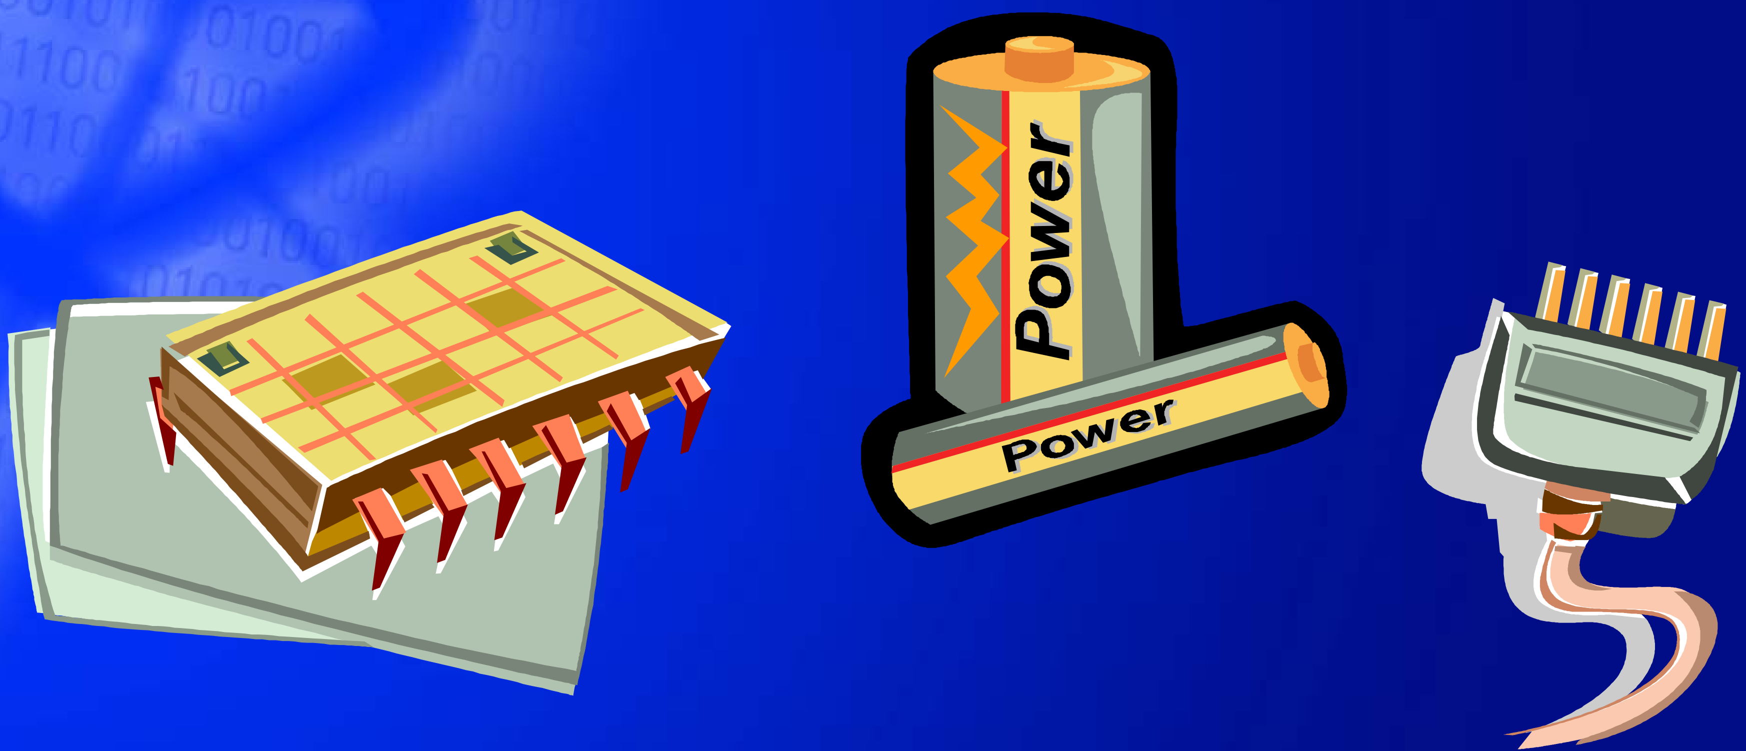

Copyright © 2002 Intel Corporation.

Dynamic Discovery, Binding, & Invocation

- Provide dynamic, ad hoc discovery, binding and invocation to services
  - Discover (possibly transient) resources and services that are available in proximity, local, and remote.
    - E.g. Location-based services.
  - Opportunistic use of local resources enhancing (e.g. CPU, storage) or sparing (e.g. power).

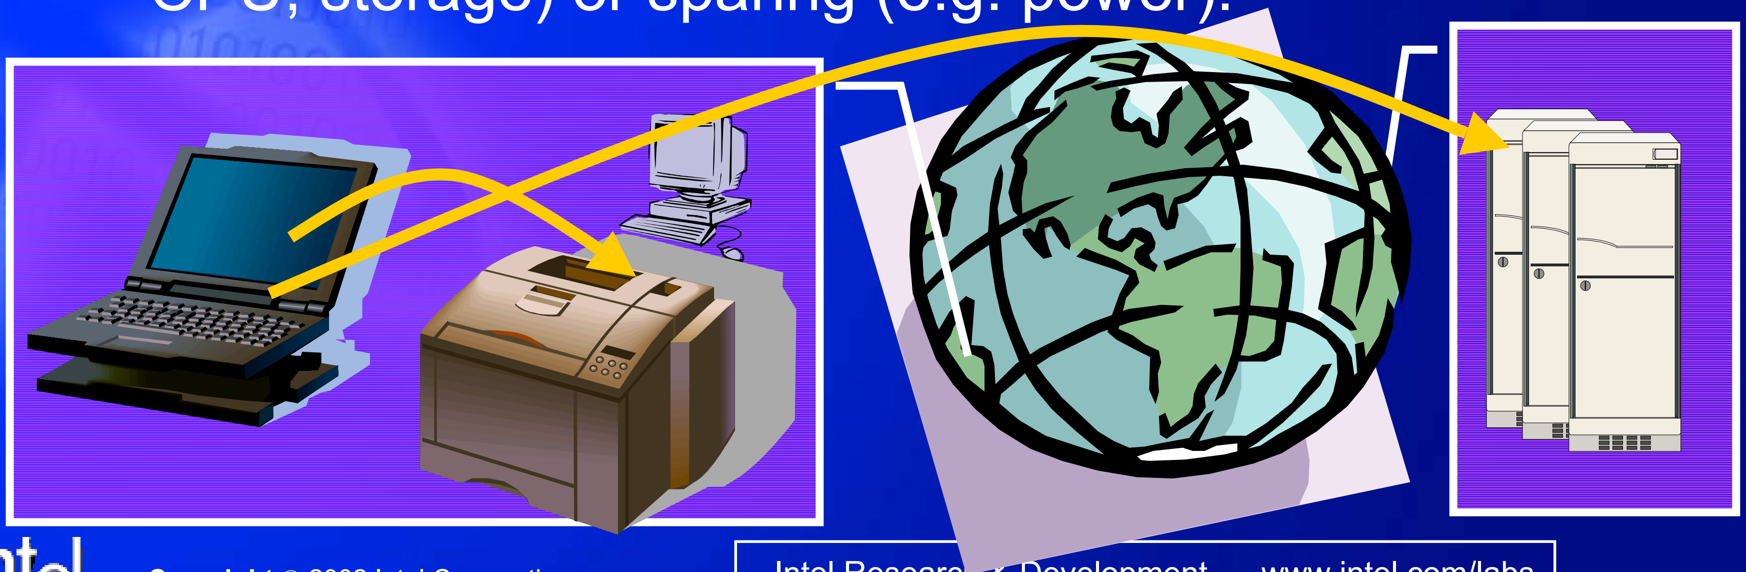

Copyright © 2002 Intel Corporation.

IDC Provides Foundation for Mobile "Magic"

- IDC with Mobility enables important capabilities for new usage models and applications:
  - Proximity and long-distance location-based services.
    - Automatic printer location, proximity synchronization with desktop.
  - Ad hoc or Spontaneous collaboration.
  - Discovery and use of compute resources
    - Cost-aware network connection sharing.
    - Cross-device resource sharing and application/data migration.

intel Copyright © 2002 Intel Corporation.

Example: Proactive Data Staging

- Mobile users accessing data not available on devices
  - Storage constraints (PDA’s, other mobile devices).
  - Data was unavailable (i.e did not exist (e-mail)) when device disconnected from network.
- Data is staged proactively for use consumption at local hotspots (airport lounges, hotels etc).
- Data can be staged on demand at hotspots (Starbucks etc).

Efficient, data access for mobile clients by moving user data closer to the user (proactively or on demand)

Copyright © 2002 Intel Corporation.

Benefits of Proactive Staging

- Reduced latency for data access for mobile clients.
- Enhanced storage capacity for constrained devices such as PDA's.
- Efficient (faster) access to user data.
  - May not be resource constrained (such as laptops) but data was not cached due to dynamic nature of content or contextual relevance. Availability of data near to a user while away from home machine.

intel
Copyright © 2002 Intel Corporation.

IDC + Mobility with Dynamic Discovery

- Provides “Innovation commons” for development of new apps
- Enables technology for next-generation distributed computing and ubiquitous computing.
  - For example, can use for gathering context from others and from environment
  - Provides information ecosystem for advanced, proactive personal services
- Facilitates dynamic use of distributed transient resources

Copyright © 2002 Intel Corporation.

Scalability and Usage Examples

- intranet/Internet
  - a company pools its workstations, servers and (off hours) office PCs into a grid to perform crash testing
- Home Area Network
  - home PCs pool their collective cycles temporarily to compress edited home movie (e.g. into DVD-R format)
- Personal (proximity) Area Network
  - user carries some of her hardware and data, augments capabilities, such as storage and UI, from the environment
    - assumes pervasive, rich embedded computer resources in the environment, initially hot spots “smart spaces”

Copyright © 2002 Intel Corporation.

Future Use Models

- Pervasive computing in an “Internet of things”
  - Data flow from sensors everywhere
- Augmented computing
  - “Smart” environments
- Context-aware usages
  - Based on location, identity, calendar, ‘history’, …
- Automated business processes
- And many more …

Copyright © 2002 Intel Corporation.

In Summary ..

- IDC is the unifying convergence of three central computing environments
- IDC ideas may be applied to add transparent mobility to today's static distributed computing
- IDC provides foundations for proactive and ubiquitous computing

Copyright © 2002 Intel Corporation.