SE3D Programme

Overview

Peter Toft
SE3D Programme Manager
HP Labs Bristol

Localised for UK English

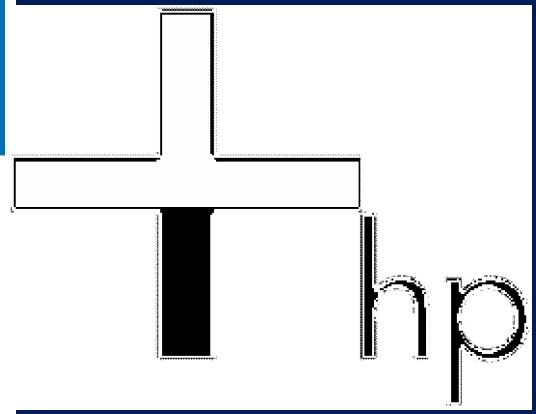

Research Technologies

- Frame Factory: a fully-featured, streamlined remote CG rendering service designed to run on a utility infrastructure, consisting of client- and server-side applications, and supporting connection over the public Internet.
- Elephant Store: a novel compression mechanism based on finding the minimum difference between one version of a data-set and the next. Used to vastly accelerate the upload of source data to the rendering service, and to store many versions of the data efficiently in the storage cache.
- Utility Security & Trust Record: protects the utility services from unauthorised access, and the service users from connecting to an impostor service. Protects customers' data from one another, and supports audit to check that the utility's automatic behaviour is compliant with expected, secure standards.
- Sumatra: the resource allocation system. Supports market-based determination of how to allocate resources amongst customers (as well as other mechanisms), and performs the actual allocations securely.
- Mercado: the user interface to the market-based resource allocation system. Presents available auctions, allows bids to be placed and credits to be managed.
- Management by Business Objectives: a service-level agreement decision technology. Decides how to re-allocate resources to customers in the event of not being able to fulfil all resource reservations (e.g. due to failures), based on criteria such as minimising penalties.
- SmartFrog: used to capture the configuration of the entire software stack as a set of templates; these drive the automatic installation, start-up and operation of thousands of software components on hundreds of machines. Also allows the software to be stopped and removed automatically.
- Anubis: A reliability protocol used to detect failures (software, hardware, network) and recover from them in a manner that is guaranteed to be timely and consistent across the complete system. Also used as a core component of Sumatra to perform resource discovery and allocation.

Utility Rendering Service: Workflow

1. Create input data in Alias Maya

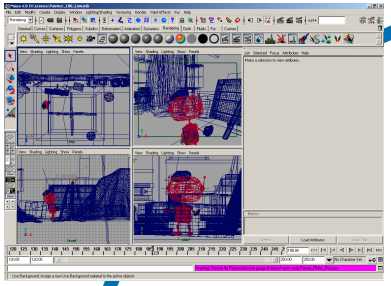

2. Upload versions of input data to the URS

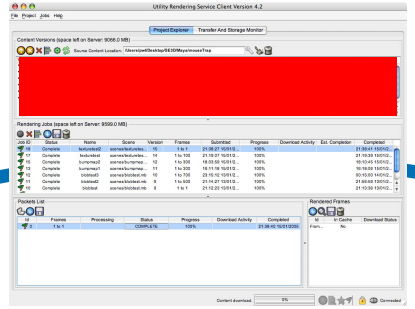

3. Manage multiple versions of input data

4. Submit rendering jobs against any version of the input data

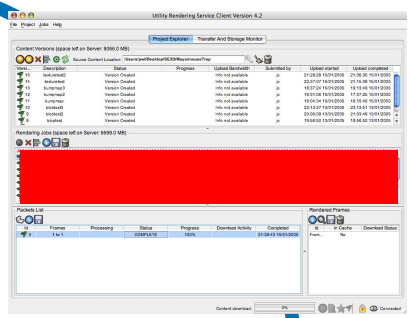

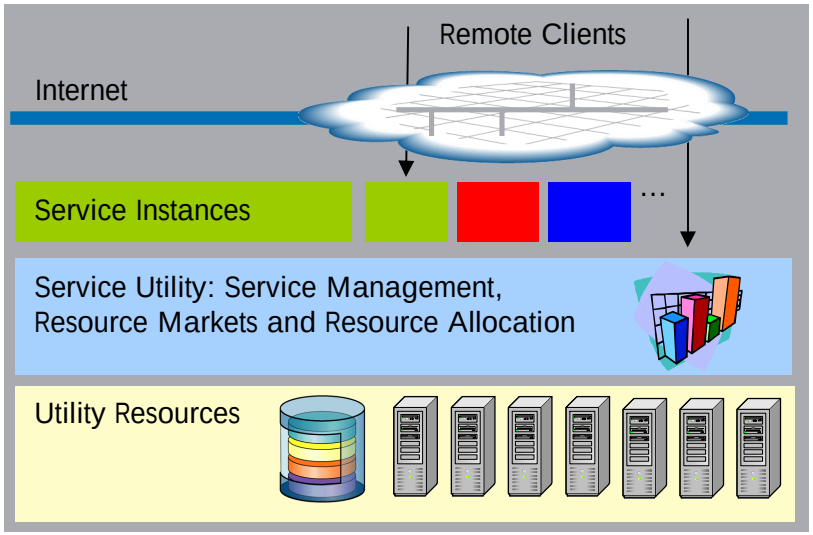

5. Purchase rendering power using the resource markets

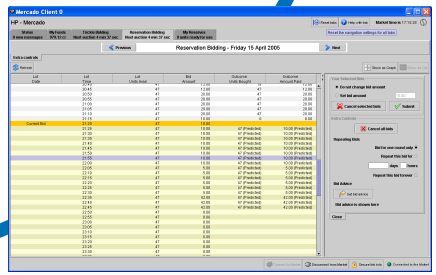

6. Manage your reserves of rendering power

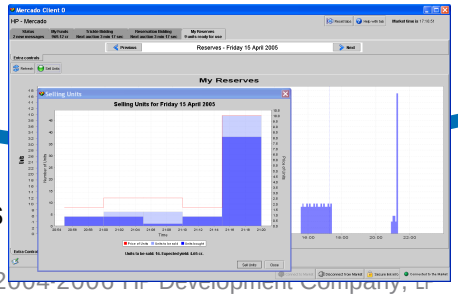

7. Download the rendered frames

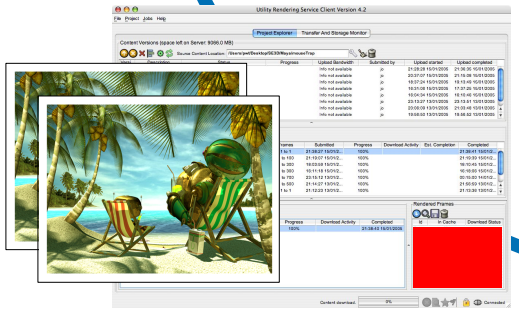

8. Repeat until your film is perfect!

Utility Rendering Service Client

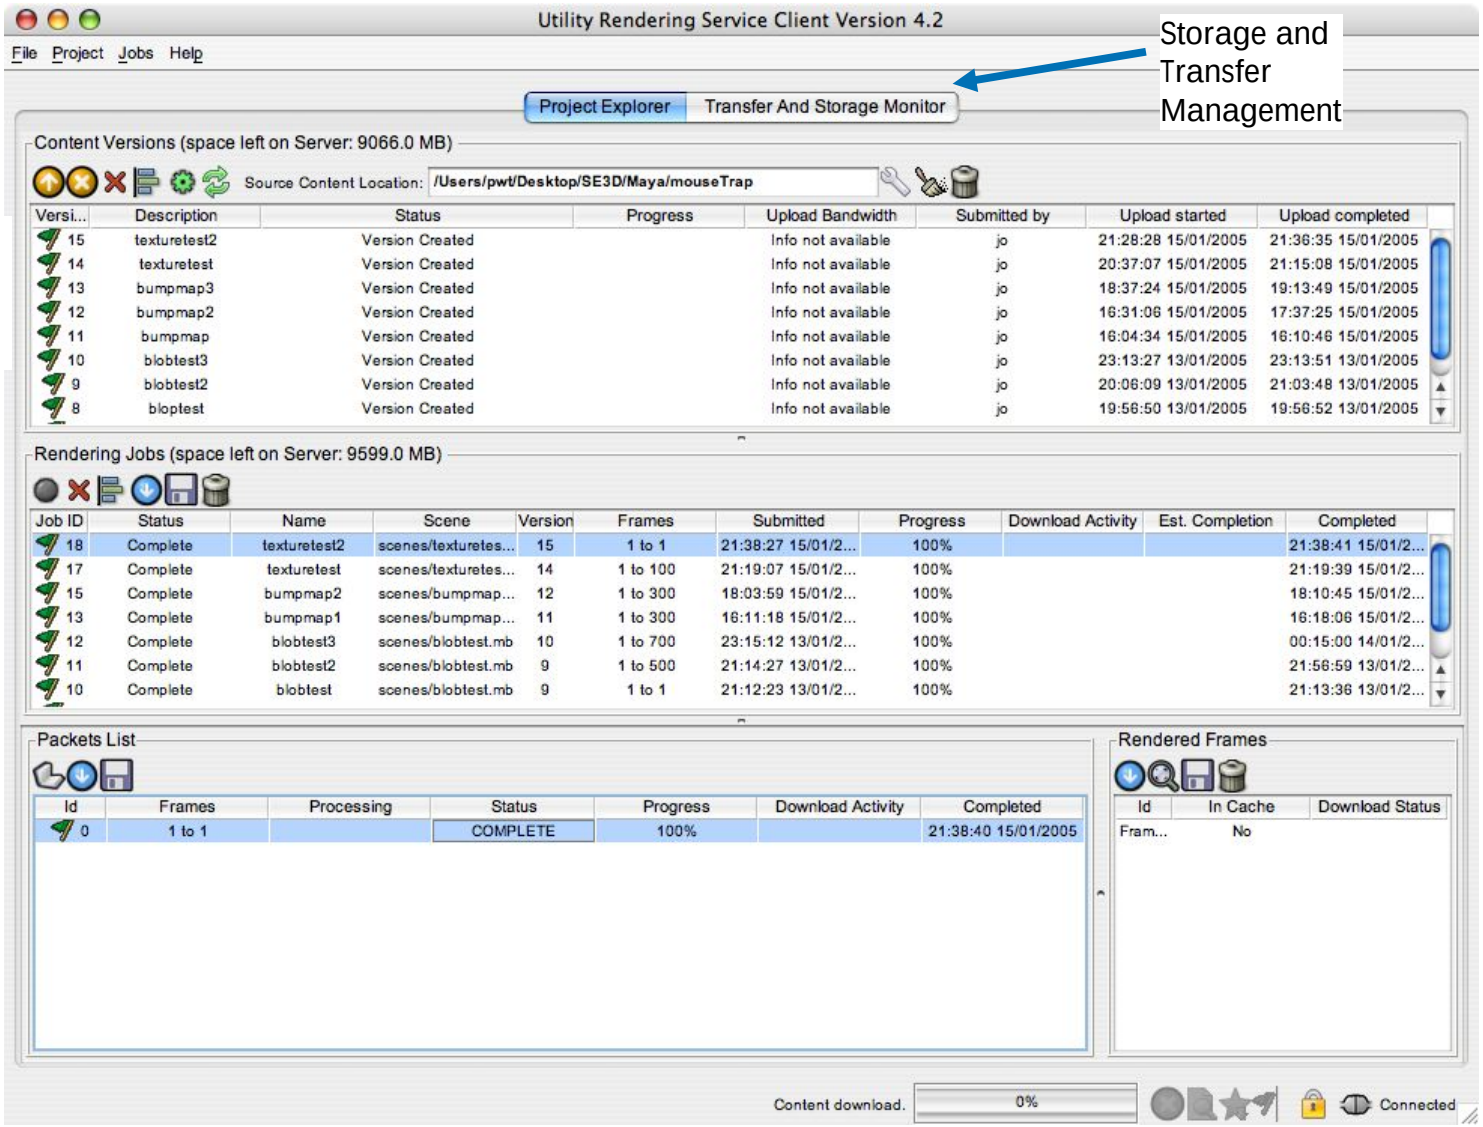

Mercado Market User Interface Example

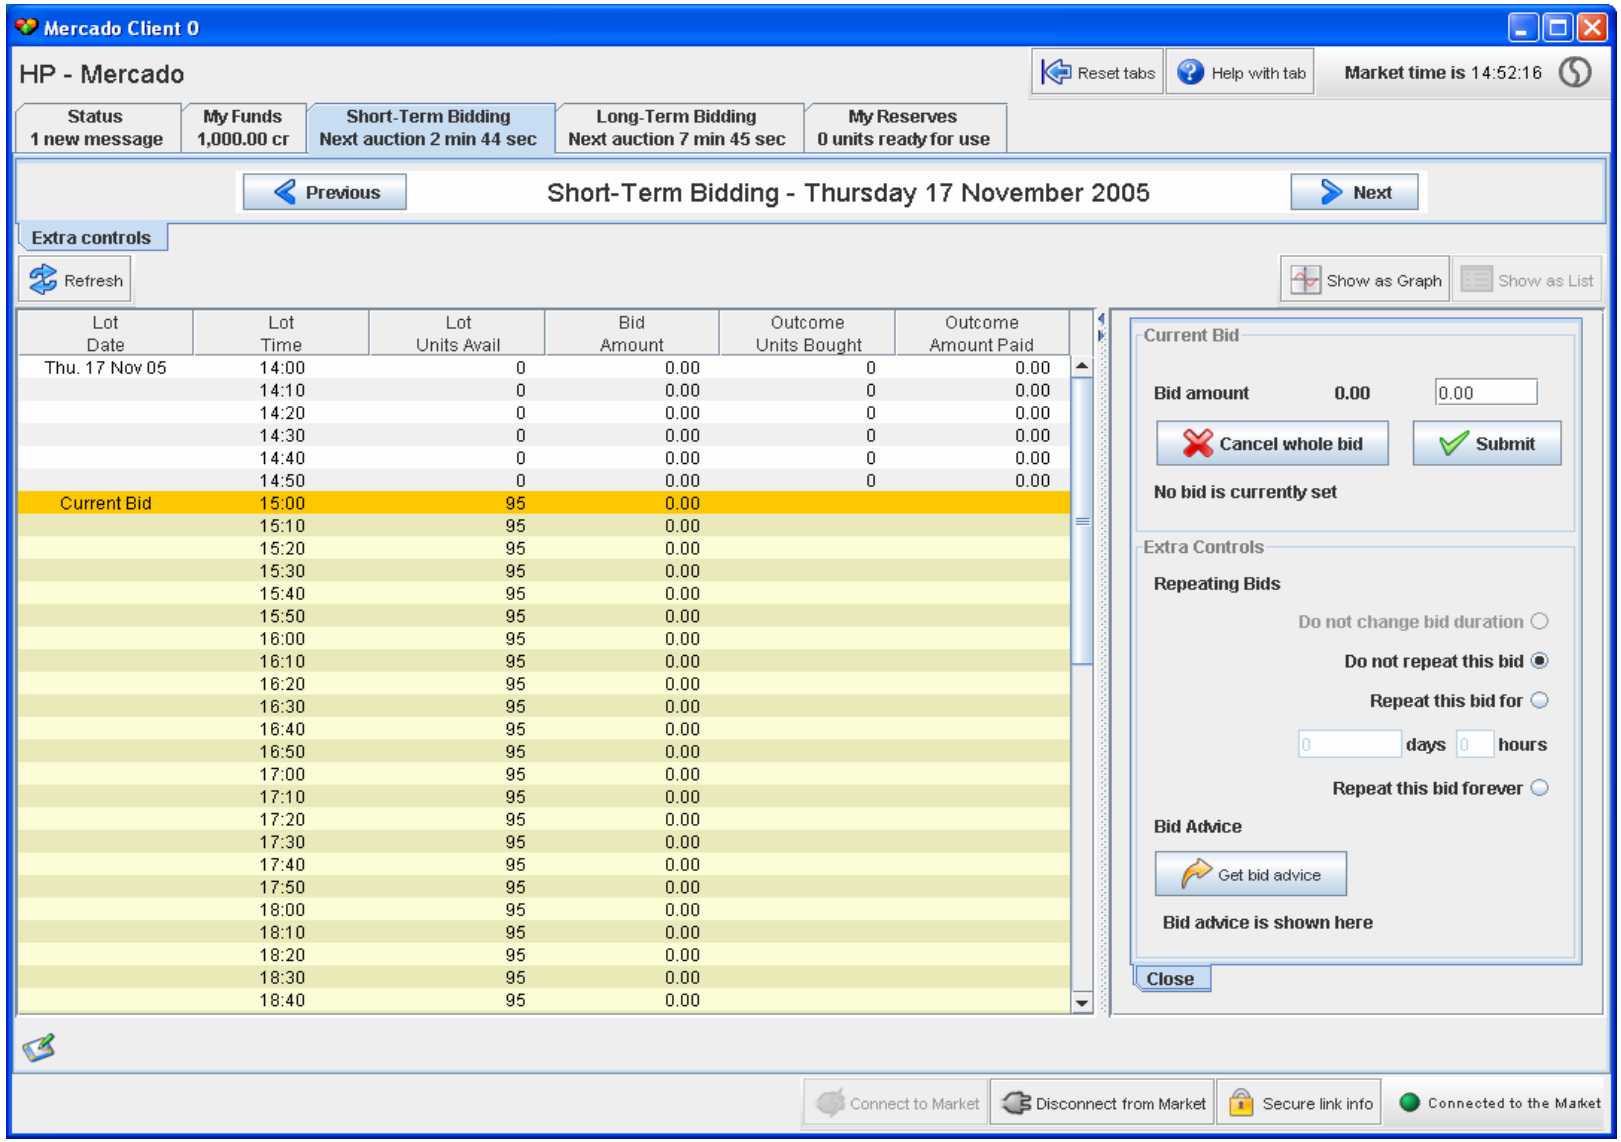

Rendering Service Architecture

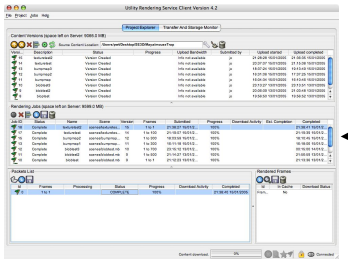

Multiple URS clients connect to separate, secure rendering service instances

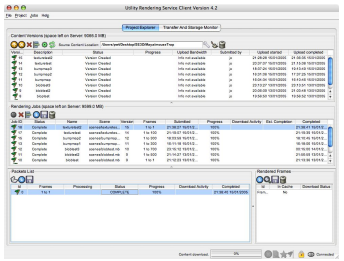

Open internet, secure transport

Optimised use of network bandwidth (Elephant Store client)

Service Utility Platform

Firewalls and Proxies

Optimised use of storage (Elephant Store Server)

Rendering Service Instance

Service Manager and Job Controller

Asset Store

Rendering Node Pool

Dynamic server allocation

Rendering Service Instance

Service Manager and Job Controller

Asset Store

Rendering Node Pool

Service Utility Platform

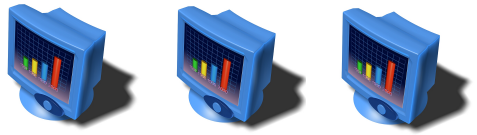

Service Instance Management

Services

Multiple Service Types and Instances

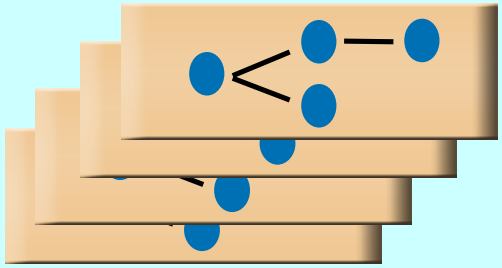

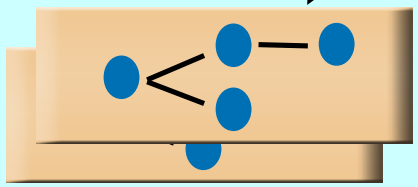

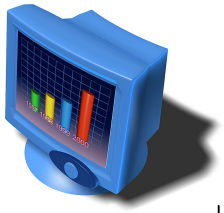

Service Utility Management

SmartFrog: Service Deployment & Monitoring

Service Utility: Maps Resources to Services

Multiple Resource Schedule Generators

- Market-based 1
- Market-based 2
- Automatic

Resource Allocation Schedules

Resource Manager / Management by Business Objectives

SmartFrog: Platform Deployment & Monitoring

Resource Pool (Proliant Essentials)

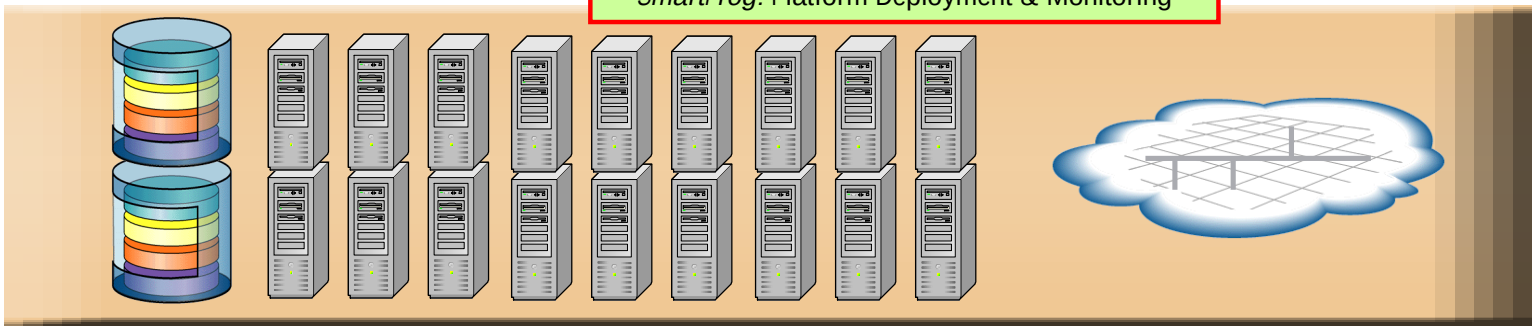

Research Results

- The Utility Rendering Service met the requirements of the SE3D film-makers
  - Worked well even over consumer-class broadband and the public Internet
  - No need to worry about configuring and maintaining the rendering servers
  - Could be a commercial service today
- Market mechanisms for acquiring computing power were easy to use
- Automation technologies performed very reliably and were key to operating the service efficiently
  - Little operator intervention required
  - Automatic resource allocation
  - Automatic failure detection and recovery
- Security measures defended against all attacks
- Some user interface changes to the URS client were recommended, e.g.
  - The ability to prioritise rendering jobs
  - A more transparent mapping between local and remote data
- For the rendering application, there is a need for a true reservation market including tentative reservations
- The market mechanisms did not experience enough contention to be tested as well as we'd hoped

Facts and Figures

- Service ran in Palo Alto, California
  - ~120 servers, 4TB Storage
- Service ran for 10 months with >99% reliability
- Peak of 15 simultaneous rendering services
- 9 films complete, 2 pending
- 500 bids placed in the resource markets, covering 5,500 hours of auctions
- 500,000 hours of CPU time ‘sold’
- 500,000 frames of animation rendered
- 1780 source data uploads
- ~30:1 compression performance on uploads
- ~ 25:1 compression performance on version storage
- More than 1,000,000 rejected connection attempts
- More than 15,000 worm attacks rejected

External Organisation

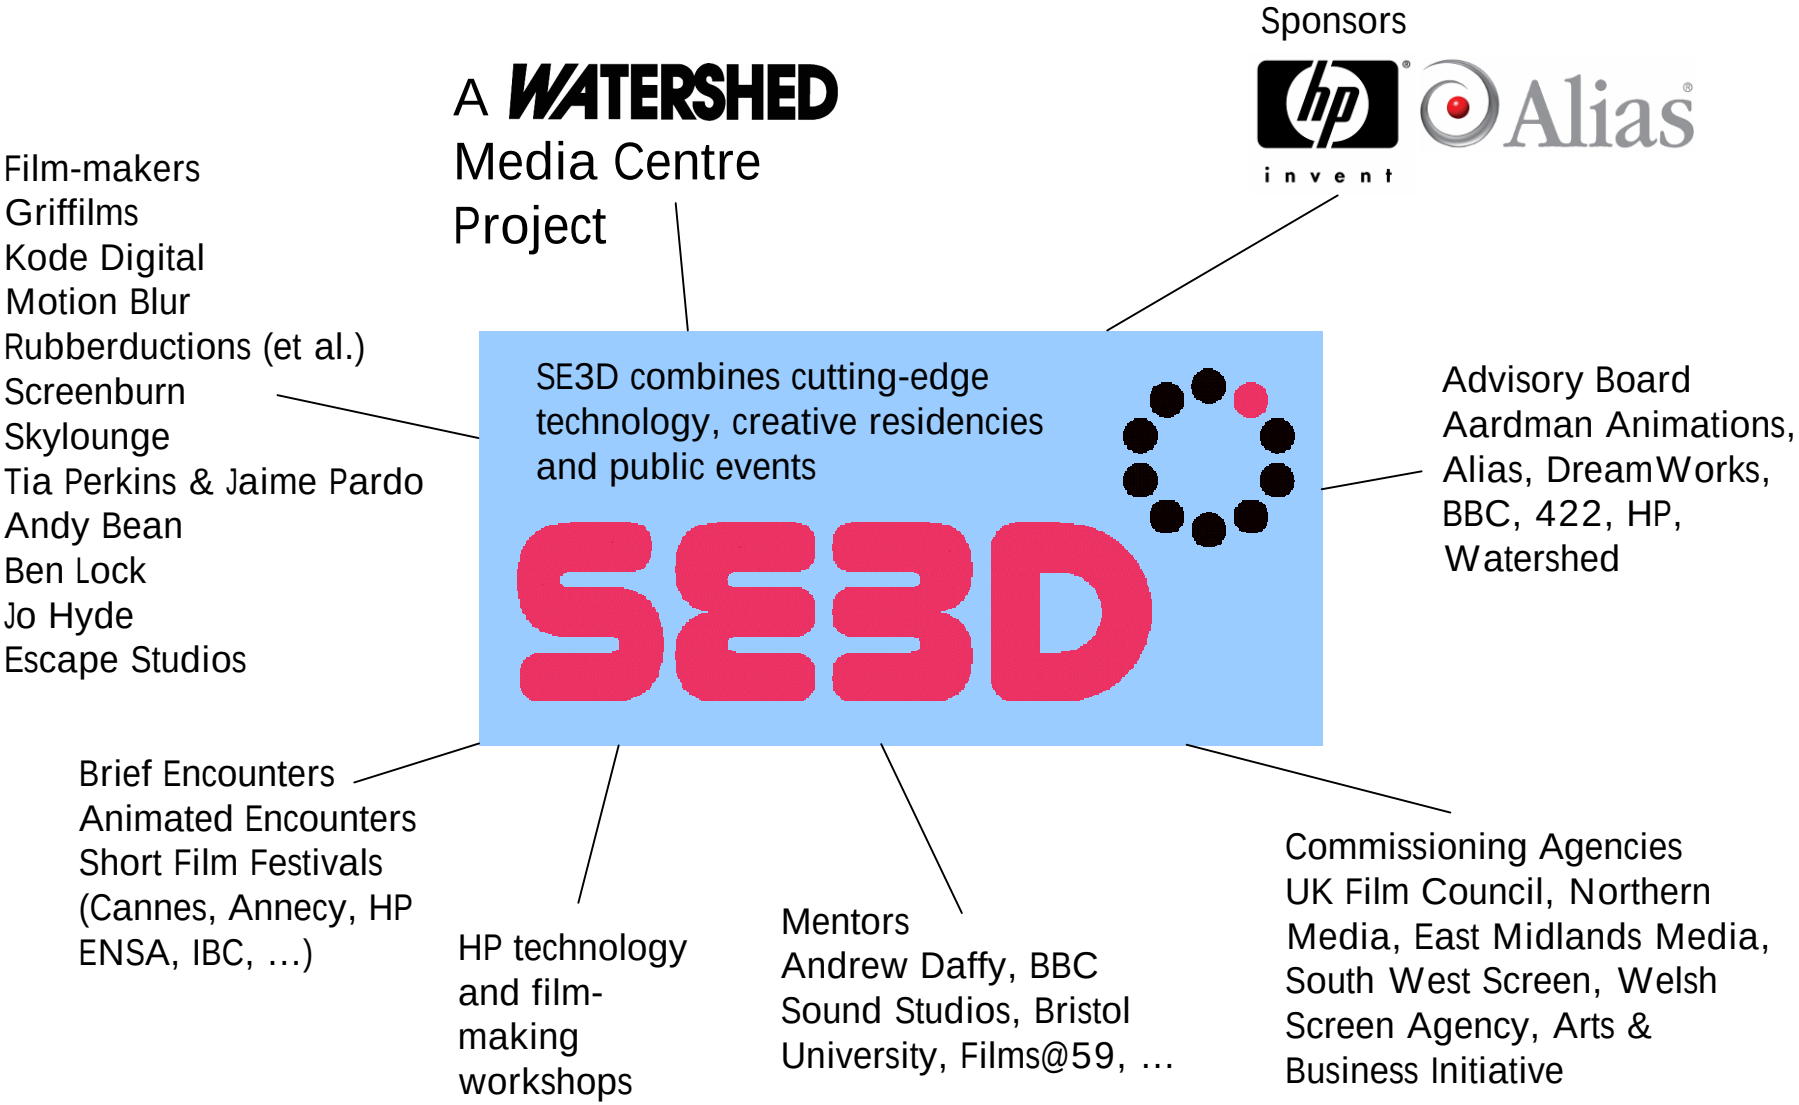


invent