

```
// innovate
public class BrainMurmurs extends Software
    implements Design, Development, Innovation {
```

An Informal Overview of Supercomputing Systems

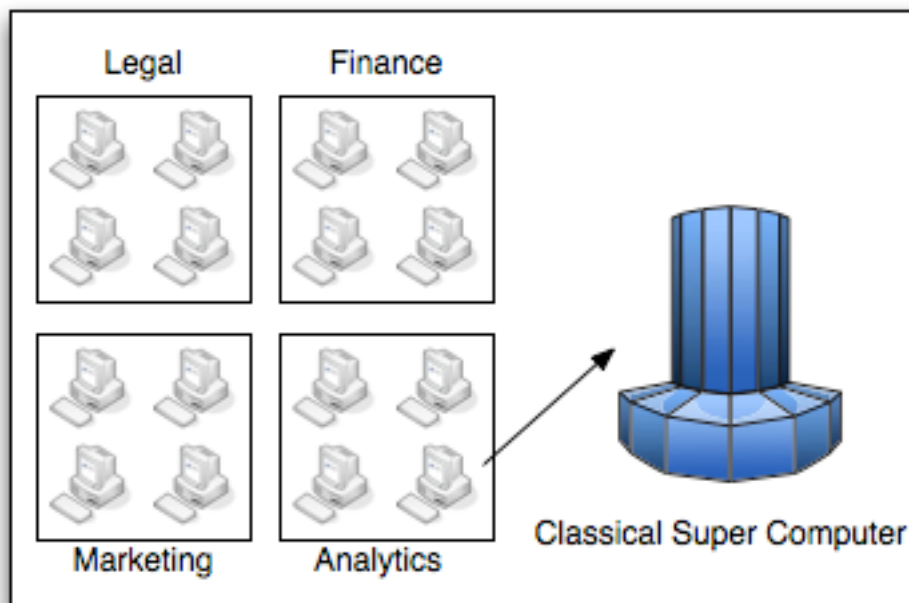
Brain Murmurs, Incorporated

This whitepaper presents a brief survey of several approaches to supercomputing. It initially discusses “dedicated hardware” approaches in which computers are purchased and housed on-site for the sole purpose of supercomputing, and then moves on to discuss grid computing systems, which involve changing network topologies, dynamic allocation of computing hardware, external resources, and computers which may primarily be used during regular business hours as office computers.

Dedicated Hardware Solutions

The Supercomputer

The original approach to supercomputing was to simply pay top dollar and buy a supercomputer. Supercomputers usually have 12 or more processors connected together by a highspeed interconnect that allows processes running on the machine to rapidly exchange information with each other. The hardware is a dedicated resource used by one group or party or timeshared between members of the organization. The supercomputer has only one purpose: addressing the computational needs of the organization.



In addition to being expensive to buy, supercomputers are expensive to own and operate. Power consumption alone can add considerably to the total cost of ownership.

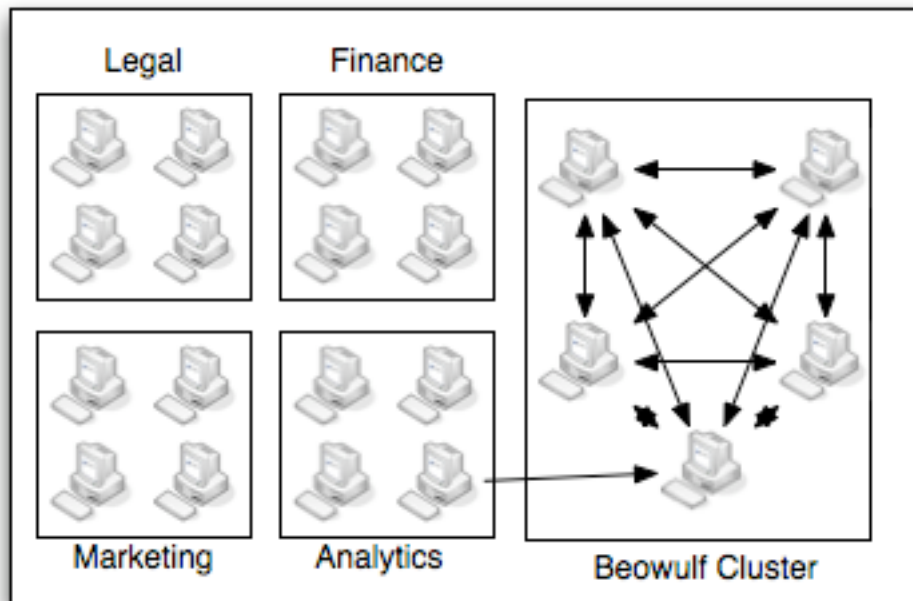
```
// innovate
public class BrainMurmurs extends Software
    implements Design, Development, Innovation {
```

Despite advances in work with commodity computers, supercomputers still reign supreme at doing jobs requiring constant communication between processes, such as Computational Fluid Dynamics (CFD).

At one time vector-based processors, such as those offered in the past by Cray, simply performed the same task in a massively parallel fashion without any communication between processes at all. This was extremely effective for handling jobs in which the same process had to be carried out on millions or billions of pieces of data.

Beowulf clusters

At some point it was noticed that the problems handled by vector-based processors could be dealt with handily on regular desktop computers. When Beowulf clusters first came on the scene, they were lauded for their phenomenal economy of cost. Beowulf clusters are networked commodity PCs, traditionally running Linux, and often connected to each other by a fairly fast network such as myrinet or gigabit ethernet.



Beowulfs were orders of magnitude cheaper than the supercomputers of the day and could work just as fast for certain classes of problems requiring fairly little or reasonably slow communication between the various computers.

In a Beowulf cluster the program to be run is distributed out to each machine from a main console via a shared file system and then run in parallel. Application APIs such as MPI allow the applications to perform different tasks based on assigned roles and to exchange data with each other if necessary. Beowulfs offer excellent performance for solving naturally parallel problems and very good performance for solving problems involving interprocess communication.

```
// innovate
public class BrainMurmurs extends Software
    implements Design, Development, Innovation {
```


Although clusters are less expensive than supercomputers they require a substantial outlay of time and money to maintain. Setting up a Beowulf is highly nontrivial, and power consumption is another factor that needs to be considered. In addition, the cluster must actually be physically housed somewhere.

Grid Computing (Time Share) Solutions

The basic notion of grid computing is simple: the actual usage of most computers is well below 100% and there are many people that have supercomputing problems but don't want to buy or can't afford to buy dedicated hardware of their own. The solution is to time-share the existing resources, either for fee or for free.

The term "grid" is a reference to the existing power grid: the computer grid is simply an available set of pooled resources that researchers can tap into on an as-needed basis, similar to how we use power from our local power company without regard to its original source (hydroelectric dam, windmill, pony, etc). The resources can be individual computers, compute clusters, or supercomputers.

Leased Grid Services



Leased Grid Services

In this case the work is clearly farmed outside the organization. Large infrastructures are being researched by the Global Grid Forum and other organizations to determine better ways to dole out work, locate available resources, bill for time borrowed from grid participants, and how to compensate participants for the use of their systems.

Grids can handle a very wide range of problem types. The cost benefits are legion, as all of the issues associated with purchasing, maintaining, administering, powering, and

```
// innovate
public class BrainMurmurs extends Software
    implements Design, Development, Innovation {
```

housing the computational infrastructure simply vanish. Sun and IBM currently offer leased commercial grid services. It costs a dollar to run your application on one of Sun's leased CPUs for an hour.

Volunteer Supercomputing


SETI@home popularized the notion of volunteer supercomputing, in which individual users and groups run screensavers and participate in a grid when the machines are not in active use.

SETI@home was able to create a vast virtual cluster by sensationalizing their application (which searched for proof of extraterrestrial intelligence by analyzing interstellar radio data) and getting hordes of people to run the program for purely altruistic needs.

The decision to embed the processing client in a screen saver was frankly brilliant, as it provided as simple mechanism for ensuring that the computer was not currently in use.

Virtual Clusters

Virtual clusters are a form of volunteer supercomputing that typically on computers world wide, without any real organizational boundaries. All data is transferred to the primary server across the internet.



Virtual Cluster

SETI@home-style virtual clusters are fantastic for handling naturally parallel applications, in which the individual computers do not need to communicate with each other. In SETI's case they had a vast pile of radio data packets to process and a standard application for analyzing a packet. The problem is considered to be naturally parallel because the radio data could be processed twice as fast by 40 computers as 20 computers (perfect scaling).

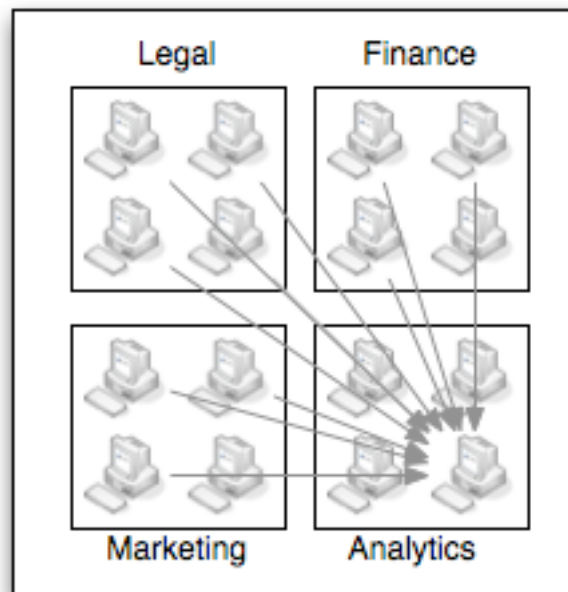
```
// innovate
public class BrainMurmurs extends Software
    implements Design, Development, Innovation {
```

There are a plethora of batch-processing jobs similar the SETI@home problem in a wide variety of industries. Possible examples include Monte Carlo simulations, computer animation, protein folding, gene sequencing, and brute force cryptography. Again, in most cases the application and data are distributed outside of the organization.

Virtual Private Clusters

Virtual private clusters are volunteer supercomputing systems contained within a single organization. In this case the applications are run on employee desktop PCs and servers within the organization on a part-time basis. This allows organizations to effectively utilize their own computers in an inexpensive and affordable manner.

Although used during the daytime for email, word processing, project management, software development, surfing, etc, these machines can all join into the computational effort when not in immediate service to their human users, i.e., at night, weekends, and perhaps during meeting times and lunch. It is imperative that the grid solution does not interrupt or affect the working environment of the staff or the business.



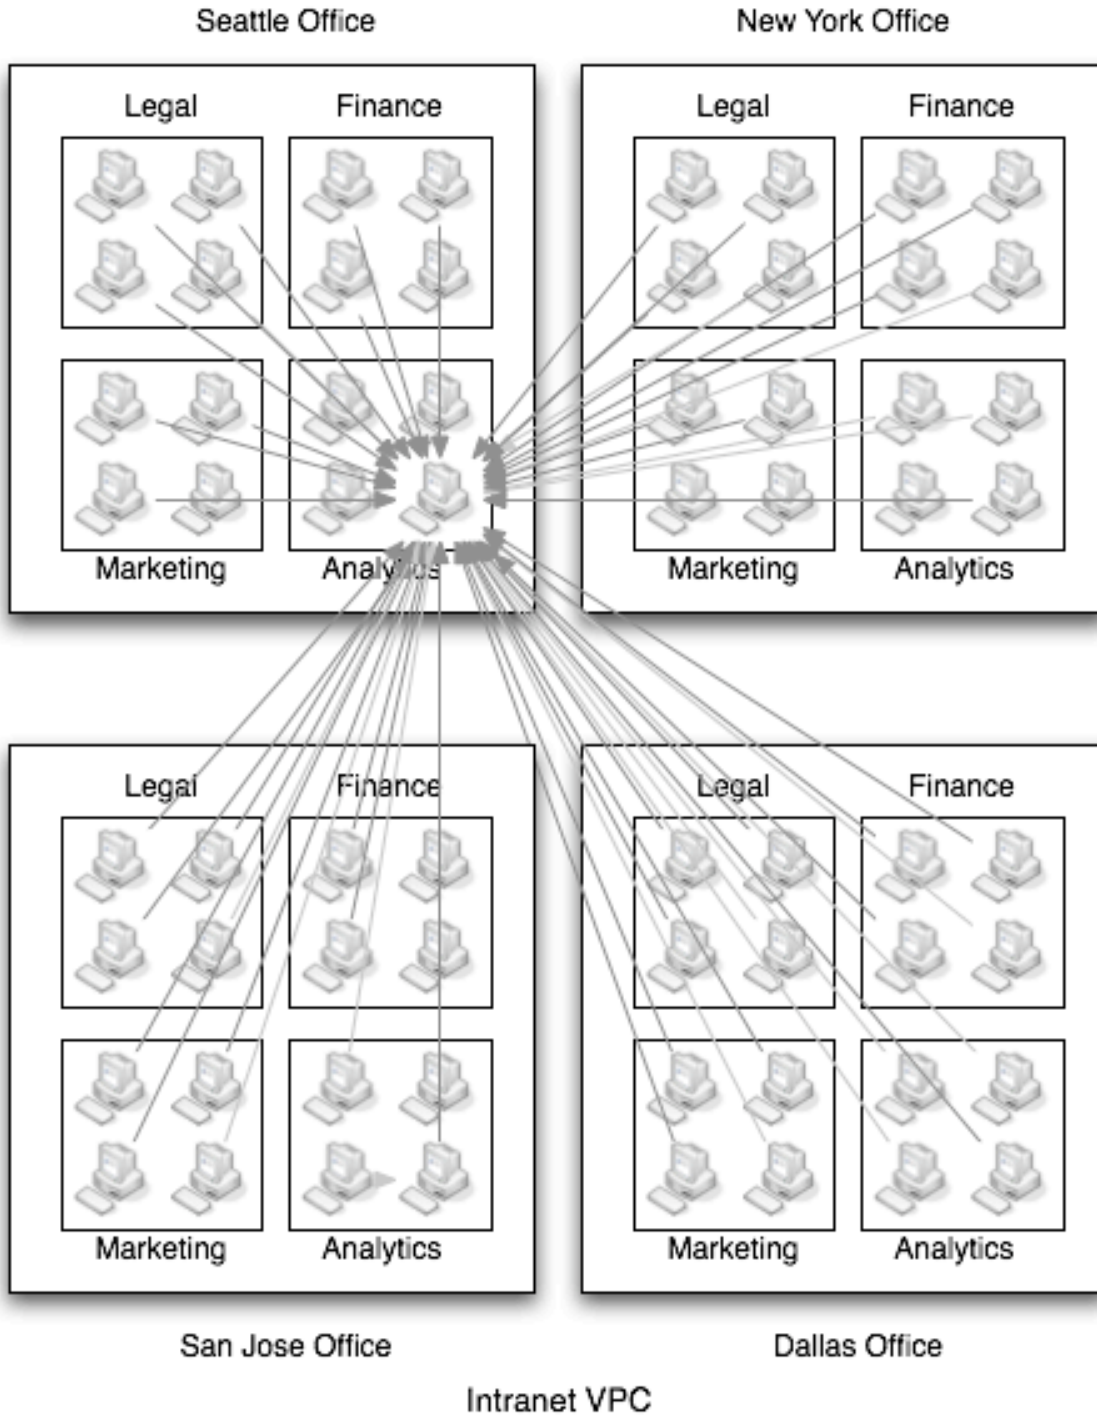
Virtual Private Cluster

The benefits to this approach are that costs are low and all data is kept inside the organization (if security is a concern). The problem with this approach is that a bug within the grid application used on the donated machines could conceivably bring down the organizations entire IT infrastructure and completely halt all business until the situation is resolved. In worst case scenarios there may be data loss as well.

```
// innovate
public class BrainMurmurs extends Software
    implements Design, Development, Innovation {
```

Virtual Private Grid

In this case, large organizations with offices around the world, such as Lockheed Martin, Boeing, and hopefully some day, Brain Murmurs, can pool their resources to create an extremely large intranet-based virtual private cluster.



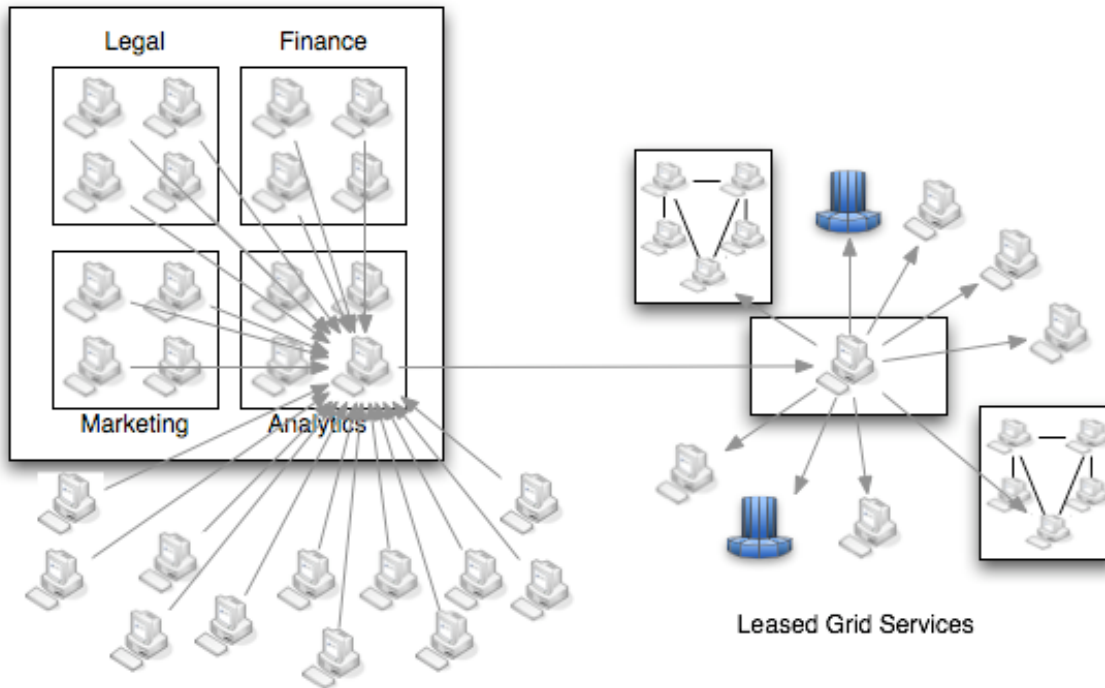
```
// innovate
public class BrainMurmurs extends Software
    implements Design, Development, Innovation {
```

Since this will involve traversing a firewall and send data across the internet an ideal solution would be to use HTTP protocol over port 80, which is typically permitted even in very secure organizations. If encryption and authentication are required the protocol should probably use HTTPS and require login credentials with client requests.

One noted advantage to this approach is that dual shore organizations (perhaps with offices in the United States and India, for instance) will have a constant supply of available grid nodes. If machines are in use in India, it is probably at night in the USA, and vice versa.

Composite Solutions

Finally, there are composite solutions, in which companies may use machines within their organization, volunteered resources from the public, as well as leased grid solutions to do supercomputing.



Composite Grid Solution

This may be appropriate in cases where every computer possibly available must be brought to bear on the problem, or in cases in which organizations prefer to defray their computing costs by running their naturally parallel applications on more affordable systems and farming out their interprocess communication-intensive applications to systems better suited to the task.