Harness Energy like a Resource
Energy Thrifty

Hi.
I'm Tony Harvey.
I'm here to talk about
HP Thermal Logic Technology.

CIO and IT staff
have 3 large problems:

• their customers want latest high performance servers;

• they want to consolidate and simplify;

• and their facilities folks say the power consumption is too high and that they don't have the air conditioning capacity to keep up.

Why is this happening?
Because Power Density is increasing.

Let's take 10KW of power in a rack and see what happens to the number of systems we could power.

Back in 2000 you could have completely filled a 42U rack with Proliant DL360 servers, but as we've moved through the generations increasing in performance and power consumption the number of servers per 10KW has dropped until here we are today at only 17 servers.

And this is the core problem it's not that Rack Mount servers or Blade Servers are difficult to power and cool.

The problem is the increase in power density.

Of course there is always the up side and that is that performance has increased as power consumption has increased meaning we can get more work done.

Of course when you increase power density you also increase heat density.

As you can see it's really just math heat output measured in British Thermal Units or BTUs is the input power in Watts x 3.41.

And for every Watt that goes into the server it takes another 0.5W – 1.5W to run the cooling the system.

• So for a system in the USA if we have 1MW of power for servers
(100 x 10KW racks)
• Takes 0.5MW to cool
• At 10cents/KWhr
• = $150hr or $1.3 million per year.

These costs are significant and reducing them directly affects the bottom line of every company that uses an HP BladeSystem.

So, how does Thermal Logic help reduce datacenter power consumption?

Well, the HP Engineering teams worked across the board from Blade design to using power efficient components from partners like Intel and AMD to even inventing some new technologies to make cooling less power hungry.

First at the most basic level we made the power sub-system more efficient.

Power supply efficiency is simply a ratio of the input power to the output power so a 50% efficient power supply would need 2000W input to get 1000W output.

Or put another way, that power supply would waste half the power supplied to it.

This diagram shows the power supply efficiency of the power supply used in the c7000.

Like any industry standard power supply as load goes up the power supply gets more efficient.

But if you look at the lower end of the graph you can see that at low loads power supplies can be very inefficient.

Looking in more detail at the graph the lower black line is the efficiency curve of a typical white box server vendor.

So the first thing that we did was use a more efficient power supply, this puts on us the blue line.

So for any output power we are more efficient.

The next thing we did was look at a system with multiple power supplies that load share.

In the c7000 we have 6 power supplies in a fully redundant configuration.

So with light loads none of those power supplies are operating at optimum efficiency.

What the HP Dynamic Power Saver feature does is monitor the power draw then optimizes the power supply efficiency by placing some of the power supplies in standby.

This moves us up to the RED curve to maintain high efficiency across a much broader range than could normally be achieved.

Let's look at an example system.

Here we have a system that's drawing 1800W.

It has 6 power supplies supplying 300W each running at 75% efficiency.

That means 2400W input power for 1800W output power.
Or put another way 600W wasted.

Running HP Dynamic Power Saver we place 4 power supplies in standby.

So now we have 2 power supplies supplying 900W each at 89% efficiency.

Now we only use 2048W input power for 1800W output.

That's a saving of 352W per enclosure.

The other 4 supplies are in standby and are available immediately if power draw increases, or to take over in the unlikely event of a failure.

For 20 enclosures at 10 cents per KW/hr that represents nearly $5000 savings, just for the enclosure power, and that doesn't include the money saved on cooling the system as well.

HP can do this because of our pooled power architecture, where all the power from the power supplies is available to any part of the system.

Another advantage this gives us is flexible redundancy configurations so you can match your redundancy to your needs rather than being forced into a one-size fits all environment.

The next area we looked at was the Blade.

(go to model)

Here as you can see we wanted a no compromise design, a system that has the same features as a 1U rack mount server but uses less power.

The Blade design was optimized for airflow so that there is no wasted air that isn’t used to cool the system.

The heatsink is designed to give superior cooling using less airflow and each component was carefully placed to ensure that heat sensitive components like disk-drives are in the front, so that they get cool air directly from the air-conditioning system, rather than being placed in the rear of the unit where they would be subjected to high-temperature air coming from the CPU and memory.

Going from the blade to the enclosure the HP engineering team came up with a new enclosure cooling architecture to optimize the efficiency still further.

Up until today, there were really only two ways to cool blades.

Local cooling where each Blade had its own fans and was responsible
for its own cooling.

This architecture is tremendously scaleable, but uses lots of fans.

Centralized coolingon the other hand uses a single or pair of large blowers in the enclosure.

There are several issues with this design.

Limited scaleability,
you only have two blowers and once you’ve reached the limit of those blowers you have nowhere to go except expensive upgrades.

Limited redundancy, with only two blowers you can only have a single failure in the system.

And power consumption, with just two blowers the worst case device always controls the speed and therefore power consumption of the fans.

As an example if we place a device that requires 10 cubic feet per minute of airflow and one requiring 30 cubic feet per minute into the same enclosure the blowers would spin up to supply 30 CFM to every device.

PARSEC stands for Parallel Redundant Scalable Enclosure Cooling and what it gives us is an architecture that combines the best of both worlds.

Airflow through the enclosure is managed to ensure that every device gets cool air and doesn’t sit in the hot exhaust air of another device and to ensure that air only goes where it is needed for cooling.

(back to enclosure)

The interconnect modules, even though as you can see they are at the rear of the unit get cool air from front of the system through these ducts which direct the flow to rear, into the switches which the fans then pull into the central air plenum before exhausting it at the rear of the system.

The fans have a similar system using louvers that will automatically open when a fan is installed and spinning and will close again if the fan is removed or stops spinning.

By using multiple fans as you can see here at the rear of the unit we get scaleable cooling capability that can grow from 4 fans to 10 fans as your needs change.

By using a zoned cooling system, where the fans in any given zone provide cooling for that zone and redundant cooling for others means that we can provide high airflow only where it's needed...

which reduces power consumption,
uses less air and
helps to keep
the noise level down.

To reduce power consumption even further we actually add more fans, which seems to be a paradox but really works.

Looking at this slide
the blue line is the 6 fan power curve you can see that except at very low flow rates 8 fans is almost always more power efficient and as the flow rate increases 10 fans are more efficient still.

This is because using more fans allows the fans to spin slower to move the same volume of air,
so each fan
uses less power.

As an added bonus spinning fans slower means less noise,
so you get a better work environment as well as saving power.

The last piece of the cooling system is the fan.

Now a fan might seem like a simple, mundane piece of technology but to a thermal architect they are the most critical component of the whole system.

To deliver PARSEC HP engineers needed to design a new type of fan that would deliver both high airflow and high pressure in a small form factor that can scale with future cooling needs.

Let me show you
(chassis at 11K).

This simulation
is the power,
airflow and sound of 2 typical full featured
pizza-box 1U rack mount servers.

The HP Active Cool Fan drops it to this
(chassis at 6K):
Up to 66% fan power reduction,
40% less airflow required and 50% less noise.

But for investment protection, the Active Cool Fan can go up to 4 times the capability of the next best available fan,
moving at 166MPH,
to air cool a 50KW rack
if we ever need
to do this in future.

We completely redesigned the fan from head to toe so it is 100 times more reliable than traditional rack mount server fans.

But it also scales to realize every one of these benefits from 1 fan to 10 fans so that the PARSEC architecture delivers centralized cooling done right.

So -what’s the Bottom line, what does Thermal Logic from HP deliver to you?

Based on tests in HP Labs we are seeing power savings of 28% per server when compared to a 1U server.

The c7000 uses 40% less air for cooling than 1U servers.

So what Thermal Logic means to you is money saved.

On power,
On cooling
And even on buying headphones for your operators.